Report on the Research Programme of the Linear Collider Flavour Identification Collaboration

April 2007

Y Banda[7], C Bowdery[4], C Buttar[3], P Clarke[2], D Cussans[1], CJS Damerell[8], N De Groot[6], E Devetak[7,8], J Fopma[7], B Foster[7], R Gao[7], AR Gillman[8], J Goldstein[1], T Greenshaw[5], M Grimes[1], K Harder[8], B Hawes[7], H Heath[1], S Hillert[7], B Jeffery[7], EL Johnson[8], N Kundu[7], V Martin[2], P Murray[8], A Nichols[8], A Nomerotski[7], V O'Shea[3], C Parkes[3], C Perry[7], T Woolliscroft[5], A Sopczak[4], KD Stefanov[8], SL Thomas[8], T Tikkanen[5], M Tyndel[8], J Velthuis[1], T Wijnen[6], S Worm[8], S Yang[7], Z Zhang[8].

[1]Bristol University, [2]Edinburgh University, [3]Glasgow University, [4]Lancaster University, [5]Liverpool University, [6]Nijmegen University, [7]Oxford University, [8]Rutherford Appleton Laboratory

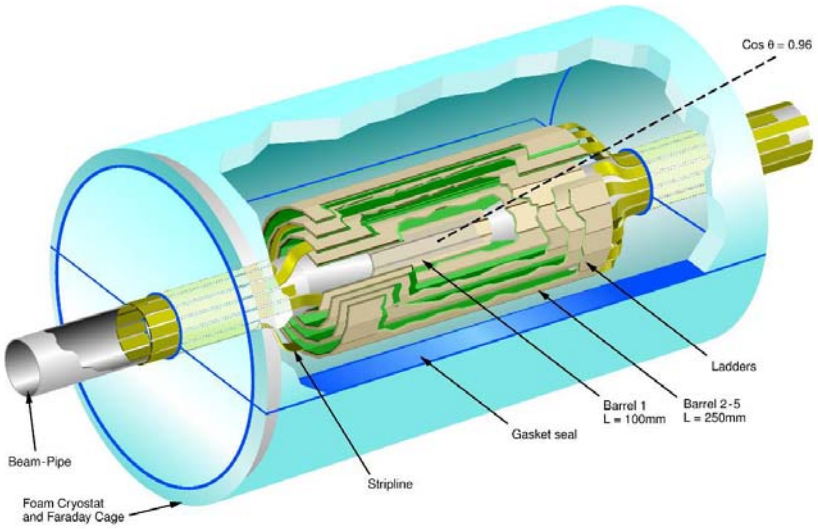

Abstract

This document describes the progress of the Linear Collider Flavour Identification (LCFI) Collaboration since the last report to the DESY Physics Research Committee in April 2005. In the intervening period, the Collaboration has successfully operated Column Parallel Charge-Coupled Devices (CPCCDs) at 45 MHz, close to the frequency of 50 MHz needed for the inner layer of the Vertex Detector (VXD) of the International Linear Collider (ILC), and has designed and manufactured a chip capable of producing the clock signals needed to drive the CPCCD. Progress has also been made with the development of Column Parallel Readout (CPR) chips for the CPCCD, allowing thresholds to be applied to the CPCCD data and on-chip clustering and sparsification. CPR chips have been successfully bump-bonding to CPCCD sensors and used for read out. The functioning of the In-situ Storage Image Sensor (ISIS) concept has been demonstrated and a second generation ISIS design is nearing completion. The Collaboration has identified materials which allow the construction of ladders with a thickness of $0.1\%X_0$, the target value for the ILC, and has constructed and surveyed test ladders using these materials. Sophisticated software for vertex finding, flavour identification and quark charge separation has also been developed and is in the process of being released. This will allow LCFI and other groups to investigate the physics made accessible at the ILC using the VXD including all experimental effects simulated in the detector concept's Monte Carlo and reconstruction packages.

Introduction

This document briefly describes the progress made by the Linear Collider Flavour Identification (LCFI) Collaboration in the two years since it last presented its results and proposed research programme[1] to the DESY Physics Research Committee (PRC). LCFI is developing the sensors, electronic systems and mechanical support structures necessary for the construction of a high performance vertex detector (VXD) at the e^+e^- International Linear Collider (ILC) and is also investigating the contribution such a VXD can make to the physics accessible at the ILC. The Collaboration's goal is to produce and test full-scale sensors with their accompanying electronics, support and cooling systems. These must allow polar angle coverage in the range $|\cos\theta| < 0.96$, readout or signal storage within about 50 μs, and have a material budget of at most 0.1% X_0 for normally incident particles, providing an impact parameter resolution of ≤ 5 μm for tracks with momentum as low as 1 GeV/c. The devices under development are the Column-Parallel Charge-Coupled Device (CPCCD) and the In-situ Storage Image Sensor (ISIS). Readout circuits for these sensors are also being investigated. Studies of materials allowing the construction of extremely low mass support structures are being undertaken. The Collaboration is also developing the tools necessary for the exploitation of the VXD at the ILC, in particular for heavy flavour tagging and quark charge identification.

Recent progress in these studies has been rapid. The Collaboration's sophisticated vertexing package is in the process of being released to the ILC community. The high-speed busline free version of the second generation CPCCD has been successfully read out at 45 MHz, close to the goal of 50 MHz required for the inner layer of the VXD. A first Column Parallel Drive (CPD) chip for the CPCCD has been designed and manufactured and a second generation of Column Parallel Readout (CPR) chip produced.

In April 2005, when LCFI last reported to the PRC, the Collaboration was seeking funding for the programme it then presented from the Particle Physics and Astronomy Research Council (PPARC) in the UK. We are pleased that, thanks in part to the positive comments made by the PRC, our application for funding until the end of 2008 was largely successful. Of the eight proposed Work Packages (WPs), seven were completely or partially funded, allowing LCFI to make progress in the areas of Physics Studies (WP1), Sensor Design and Production (WP2), Readout and Drive Electronics (WP3), External Electronics (WP4), Integration and Testing (WP5) and Mechanical Studies (WP6). Funding for the requested Electromagnetic Interference and Test Beam studies (then labelled WP7) was not provided. The final Work Package covers the management of the project.

The following report describes the work done in each of the LCFI Work Packages in the last two years, with the emphasis being on the most recent progress. Brief comments are also made on the proposed future programmes of these Work Packages.

WP1 – Physics Studies

Introduction

The aim of the LCFI physics studies is to provide guidance in the choice of vertex detector parameters by assessing different detector designs and to contribute to the development of the tools necessary to perform such studies as well as to strengthen the case for the ILC by demonstrating the physics made accessible by the collider and its detectors. Development of the software for vertex detector related event reconstruction, such as finding and fitting the decay vertices of heavy flavour hadrons, is the responsibility of the vertex detector R&D projects. With the development of the LCFI Vertex Package, the Collaboration is making an essential contribution to the library of analysis routines, MarlinReco, which work with the Marlin (Modular analysis and reconstruction for the Linear Collider) reconstruction software. This will allow reliable studies of the flavour identification capabilities of the ILC detectors and further investigations of the possibility that these can be extended to include quark charge determination.

With recent developments in other areas, notably tracking and particle flow, and the release of our code, MarlinReco is now reaching a level of maturity that will permit LCFI to make the transition to using a Geant4-based Monte Carlo (MC) simulation and state-of-the-art reconstruction tools. Only the use of such a full simulation will permit the dependence of vertex detector performance on design parameters to be realistically evaluated.

The focus of the LCFI physics effort since the last PRC review has been on the development of the Vertex Package, which is now being released. Our software, comprising approximately 20 000 lines of C++ code, is fully functional and has been carefully validated using input from the fast MC simulation SGV (Simulation a Grande Vitesse) and from the full Geant4-based Monte Carlo Mokka. These tests have shown that, where comparisons with previous studies can be made, the Vertex Package is performing extremely well. The LCFI studies have revealed some problems with the Marlin software that are currently being addressed.

The Vertex Package is urgently needed by the ILC community and interest in our work extends beyond Europe: for example, we are in contact with Norman Graf, the coordinator of US ILC software developments, who plans to test our package independently of MarlinReco as part of the international effort to ensure interoperability of the different ILC software frameworks.

After the first release of the Vertex Package, we will begin the transition towards use of a full MC simulation and shift focus to the investigation of benchmark physics processes, in parallel to continuing work on further developing the software tools for vertex detector-related studies.

The following sections describe the scope, implementation and validation of the Vertex Package, present other work recently carried out by the Collaboration in the physics area and give a more detailed outline of our future plans.

The LCFI Vertex Package

Scope of the Vertex Package

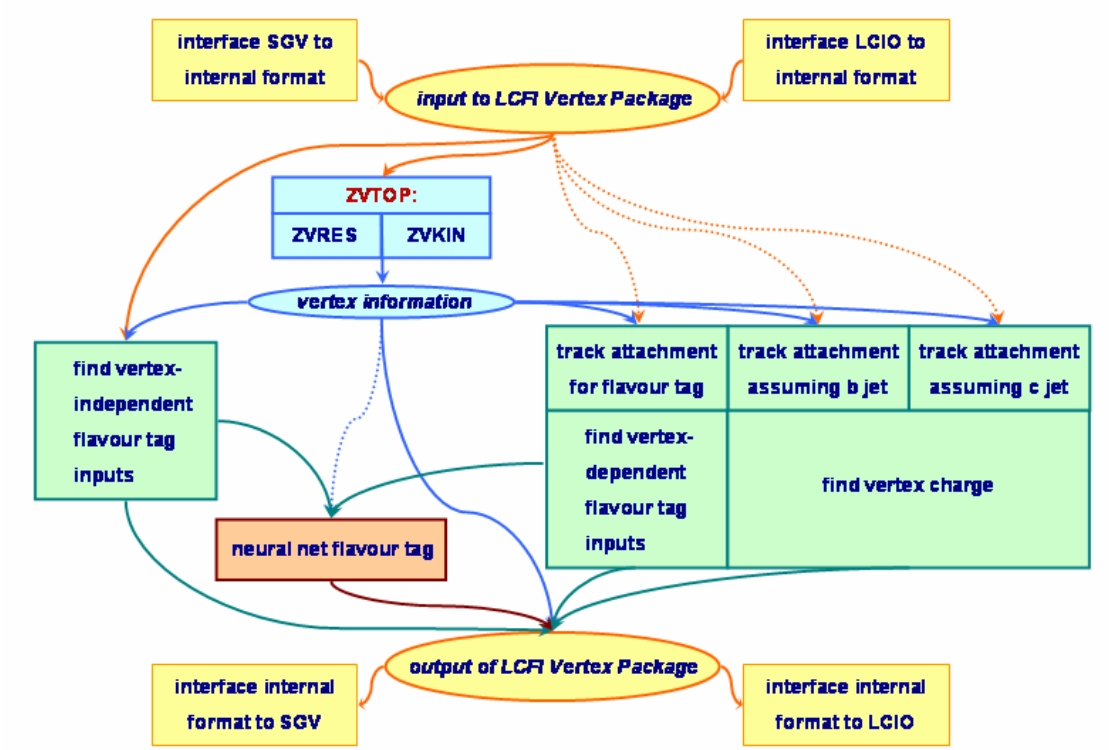

Figure 1 Structure of the LCFI Vertex Package and flow of information between the different parts of the software.

Figure 1 shows a schematic overview of the LCFI Vertex Package. Input to and output from the package is handled in LCIO[2] format, permitting data to be exchanged between the different ILC software frameworks. When used as part of the European software framework, input will usually be provided by the full MC simulation Mokka and MarlinReco event reconstruction. For test purposes, we set up an interface to the fast MC simulation SGV used by LCFI before and hence permitting detailed comparisons and cross checks with previous results.

The vertex finder ZVTOP[3] provides two vertexing algorithms, called ZVRES and ZVKIN; the latter is also known as the ghost track algorithm. While the ZVRES branch is a very general algorithm, coping with arbitrary multi-prong decay topologies, ZVKIN is more specialised. Using additional kinematic information, the ghost track algorithm is able to correctly assign tracks to vertices in decay topologies with one or more 1-prong vertices, thus extending coverage and improving the flavour tagging capabilities as shown for SLD[4], where this code was first developed. Our package is the first to provide this algorithm for ILC studies.

The default flavour tag procedure provided was developed for the ILC by Richard Hawkings[5]. It is based on nine neural networks, three each for the three cases of one, two or more vertices found by ZVTOP. Separate networks are used to identify b jets and c jets for arbitrary background. For some physics processes, the background consists only of b jets, permitting improved c jet identification. Therefore dedicated networks are provided for this case.

If the only vertex found by ZVTOP is the interaction point (IP), the flavour tag is based on information from the two tracks in the jet that have the highest impact parameter significance, plus the "joint probability", i.e. the probability for all tracks to come from the IP. If secondary vertices are found, the neural net inputs are calculated from a set of tracks which contains one or two vertices from ZVTOP, and in some cases further tracks assigned by a specific procedure devised to recover some 1-prong topologies for b-jets when running ZVRES. This "track attachment" step, in which additional tracks are added to the set of tracks contained in secondary ZVTOP vertices, may need to be tuned, depending on whether the tracks are to be used for obtaining the flavour tag inputs or the vertex charge. For vertex charge calculations, the track attachment step may differ depending on whether the jet is assumed to arise from a b or from a c quark. Track attachment for b jets was carefully optimised by LCFI in 2004 using the SGV fast MC simulation. A similar procedure will be used for c jets.

For training the neural networks and for obtaining the output from pre-trained neural nets we use C++ based neural network software developed within the Bristol LCFI group. This allows flexible definitions of the network architecture (number of layers and nodes), transfer functions and training algorithms.

Implementation of the Vertex Package

The Marlin framework, into which our code has been integrated, is based on a modular approach, permitting distributed code development. The modules, called processors, are activated and configured by a steering file. The LCFI Vertex Package provides nine Marlin processors. One processor is needed for track selection and two for the two branches of ZVTOP. In addition, there is one processor which obtains the true jet flavour from the Monte Carlo, one for fitting the event vertex, one to calculate neural net inputs and the vertex charge as well as one to train and one to use the neural networks that provide the tags. Furthermore, one processor is provided to determine and plot flavour tag purity versus efficiency. As well as these processors, the package provides steering files giving examples of how our code should be used in a typical analysis application. The output of our code makes use of a new Vertex class recently implemented by the LCIO developer's group at the request of LCFI, following discussions with other ILC groups from the US and Europe in the ILC software forum[6].

A large amount of work has been spent on the careful design, implementation and testing of the internal classes of our code that provide the equivalent of "procedures" in Fortran programming as well as permitting information to be passed around between the different parts of our code, e.g. between ZVTOP and the flavour tag. The Marlin processors access a set of algorithm classes which all inherit from a simple template class. Input to the algorithm classes are objects like jets or events. Implementing these object classes independently of existing classes, e.g. in LCIO, had the advantage of permitting them to be tailored to the requirements of our package. This advantage outweighs the fact that a dedicated interface is required between the internal data representation and the LCIO format used for communication with other parts of the analysis framework within which the package is used.

Some effort was needed to reduce the run time of the ZVRES branch of our code, which initially required about 3 seconds per jet. Systematic studies using the profiling

tool Valgrind[7] were performed, resulting in a large decrease in the run time. This was achieved by introducing caching and other improvements. Also, the absence of memory leaks, double frees etc. was ensured by the use of the automatic memory management techniques and test runs in the virtual machine environment provided by Valgrind. The resulting distribution of execution time per jet, obtained with a 2.4 GHz Pentium4 processor, is shown in Figure 2(a) and is on average within a factor of about 20 of the corresponding Fortran value Figure 2(b) shows that run time increases exponentially with the number of tracks in the input jet. Current performance is limited by the vertex fitter. Further improvements in this area may be possible in the future, but are not planned at present given the time constraints and the fact that changes to the fitter algorithm could affect many areas of the code and would hence imply extensive cross checks of performance.

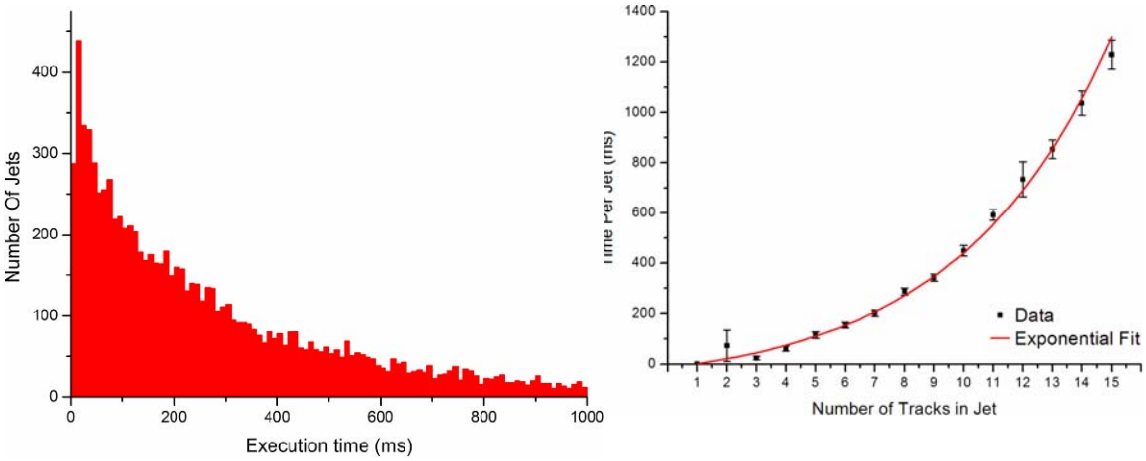

(a) (b)

Figure 2 Run time performance of the ZVRES vertex finder: (a) shows the distribution of the execution time when running the ZVRES code on a 2.4 GHz Pentium 4 processor; (b) shows the exponential dependence of run time on the number of tracks in the input jet.

For the documentation of the software, the 'Doxygen' system[8] has been used. In addition to an overview of the package and a tutorial section outlining usage of the various processors, detailed documentation of the C++ classes is provided. Development, validation and documentation have progressed in parallel.

Following the recent successful system test of the full package, the code is now in the process of being released. Once this process is complete, the package will be available via the CVS repository for ILC software maintained by DESY Zeuthen. In addition, we have agreed with the DESY group that another directory will be set up for the description of neural networks and information on the training conditions, such as physics process, sample size, input variables etc. Having a central repository for these networks will contribute to making analyses based on these flavour tags more transparent and should facilitate comparisons between similar studies performed by different ILC groups.

The release of the package will be announced at the ECFA simulation workshop in Orsay on the 2nd...4th May 2007. For those interested, tutorials on using the package will be run in addition to a presentation on functionality and validation results. Initial test results obtained with our code will also be presented at the Linear Collider workshop at DESY.

Validation of the Vertex Package

Input from the fast MC program SGV was used for the first tests of our code. We used the Fortran-to-LCIO interface available on the LCIO web site to write out events generated and reconstructed with SGV in LCIO format. This LCIO file was then read in by the Marlin processors under test, which created another LCIO file with the output of our software. For the parts of the package that have equivalents in the Fortran code, this approach permitted direct comparison with the Fortran version running on identical input events.

Starting from simple checks at the single track level, increasingly complex aspects of the code were tested, leading up to tests of the full ZVRES branch within the SGV framework.

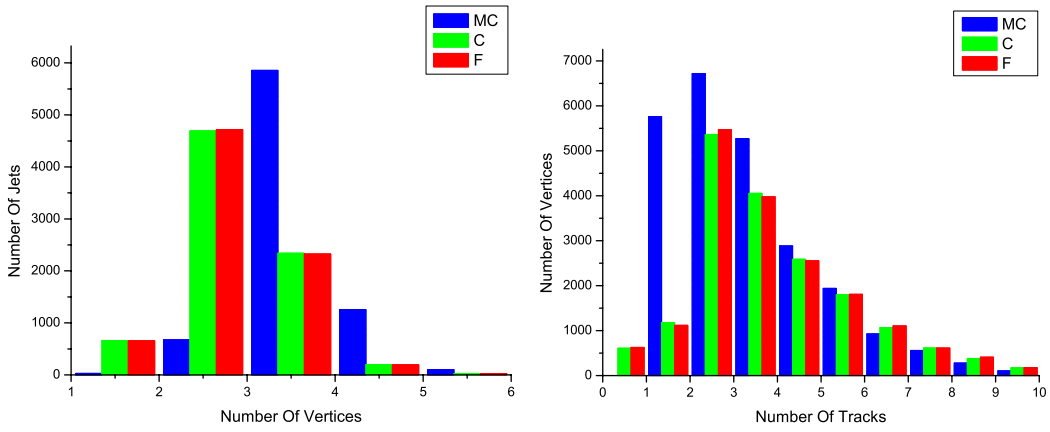

Figure 3 Distributions of number of vertices (left) and of tracks per vertex (right) for C++ (C) and Fortran (F) ZVRES and for the MC, obtained from a sample of 100 GeV b jets.

Initial tests focused on identifying and correcting various errors in track swimming and vertex fitting. To this end, the IP and one track were input to the code, with errors set by hand, to allow checks of whether the track is identified as part of the IP up to the expected threshold value of the distance between the two and that, beyond this distance, the two objects are resolved. Similar tests were later performed feeding two tracks plus the IP into the code. For these conditions, agreement between the Fortran and C++ versions is now at the level of 99.8%, the remaining 0.2% being due to differences in the implementation of the vertex fitting in the two versions.

C++ ZVRES							
Monte Carlo track origin	Two vertex case			Three vertex case			
	pri	sec	iso	Pri	Sec	ter	iso
Primary	98.1	0.7	1.2	96.9	2.2	0.1	0.9
B decay	8.3	75.6	16.1	2.3	89.1	4.7	3.9
D decay	2.3	79.4	18.3	0.6	17.5	77.5	4.5

FORTRAN ZVRES							
Monte Carlo track origin	Two vertex case			Three vertex case			
	pri	sec	iso	pri	sec	ter	iso
Primary	98.0	0.6	1.4	97.3	1.5	0.0	1.2
B decay	9.3	74.8	15.9	2.6	90.6	3.9	2.9
D decay	2.5	81.1	16.4	0.6	17.2	79.0	3.2

Table 1 Purity of reconstructed track-vertex association (%) for the C++ and the Fortran ZVRES branch of ZVTOP for 100 GeV b jets. Abbreviations for reconstruction level track assignment are “pri” (primary), “sec” (secondary), “ter” (tertiary) and “iso” (isolated).

Using a sample of 100 GeV b jets, Figure 3 shows the agreement between the C++ and Fortran versions of ZVRES in the number of vertices found and in the number of tracks per vertex. Differences between reconstructed number of vertices and the MC distribution result from the complexity of the decay chain, in which distances between the B hadron decay and the IP or between D and B hadron decays are frequently too short to be resolved by the detector.

Track-to-vertex association, in the form shown in Table 1 of the original ZVTOP paper[9], was studied for jet energies ranging from 25 to 250 GeV. Table 1 here gives an example for 100 GeV jets. The first row of this Table lists the percentages of tracks coming from the IP at the “MC truth” level which are assigned to the primary, secondary and – if three vertices are found – tertiary vertex, or are left over as isolated tracks. The following rows show the equivalent numbers for the secondary and tertiary MC vertices. Excellent agreement of the values between the two versions of the code is seen. This agreement is similarly good over the full energy range studied. The fraction of jets with exact agreement in the number of vertices found and of all tracks assigned to them is between 70 and 75%, depending on the jet energy.

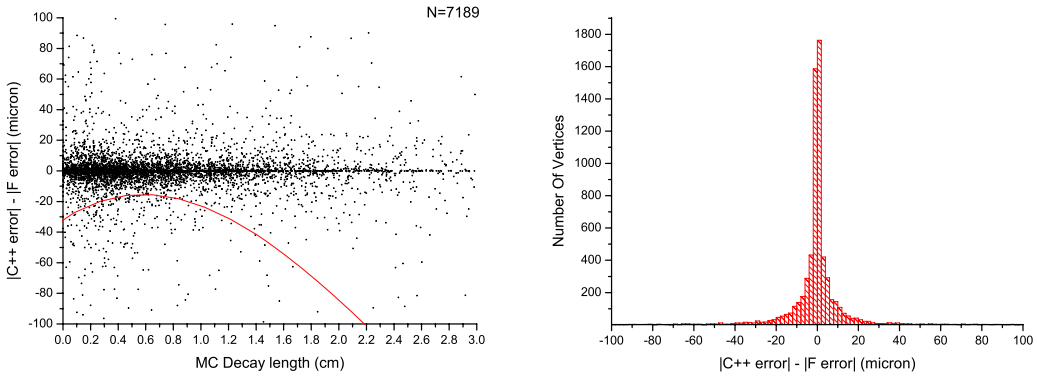

(a) (b)

Figure 4 Difference between the C++ and the Fortran deviation from the MC decay length, plotted as a function of decay length (a) and projected onto the y axis (b). The red line in (a) shows a second order polynomial fit to the distribution.

Physics performance not only depends on the track content of the decay vertices but also on how well their position is obtained from the fit to that set of tracks. This can be measured by calculating the difference between the reconstructed decay length and

the MC value. In Figure 4(a), the difference found from the Fortran version of ZVRES is subtracted from the difference found for the C++ version and plotted as function of the MC decay length. Thus values smaller than zero correspond to cases in which the C++ comes closer to the MC truth, that is, performs better than the Fortran. For most vertices, agreement is very good. In some cases vertices are only found by the C++ version. Figure 4(b) gives a projection of the plot shown in (a) onto the y axis, with the peak around zero corresponding to small differences arising from the vertex fit and the tails being due to differences in vertex finding. The red curve in part (a) of the Figure shows a second order polynomial fit to all the points, indicating that, on average, the C++ version ZVRES now performs better than the Fortran version.

After the C++ code had been integrated into Marlin, the calculation of each of the input variables for the flavour tag neural nets and of the vertex charge was checked by comparing the one-dimensional distribution from the Marlin processor to that from the Fortran code and by plotting the values from one versus those from the other. Figure 5 shows the 2D-plot for the vertex momentum. Excellent agreement is found here and for the other input variables. It is expected that a few points do not exactly lie on the diagonal of the two-dimensional plot, since there are some jets for which results from the vertexing step differ. From the previous tests it is known that in these cases, the Marlin version on average performs better than the Fortran.

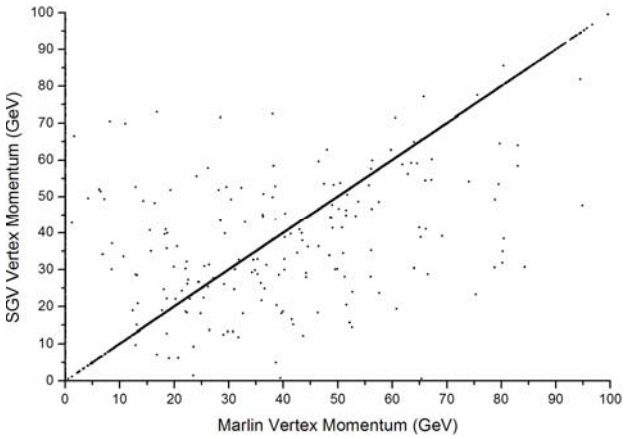

Figure 5 The momentum of the vertex furthest away from the IP is plotted. The Fortran value is plotted against the MARLIN one.

A system test of our package was performed by running the full chain of the ZVRES branch for vertex finding followed by determination of the vertex charge and by calculation of the flavour tag inputs, which were fed into the neural network processor. The resulting LCIO output file was converted into a root ntuple for test purposes and read in by the analysis software used previously, to obtain purity versus efficiency graphs corresponding to the resulting flavour tag.

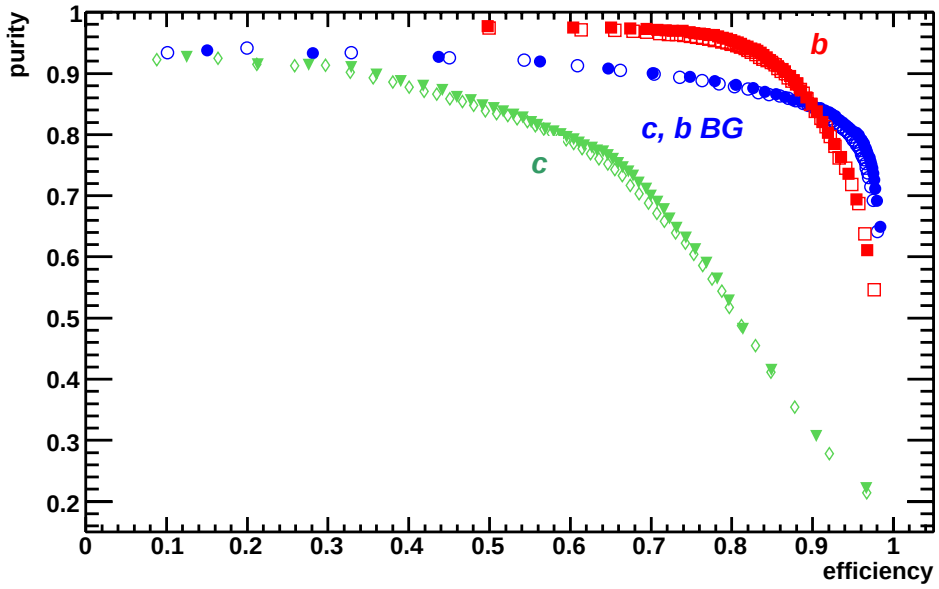

Figure 6 The purity of the flavour tags for b jets, c jets and c jets with b background only, plotted as a function of efficiency, for a centre-of-mass energy of 92 GeV. Full symbols correspond to the result from our package, embedded into Marlin and fed with input from SGV, open symbols to the performance of the Fortran code.

These graphs are shown in Figure 6 together with the Fortran result obtained for identical input events. Excellent agreement between our code and the Fortran version is seen in all cases.

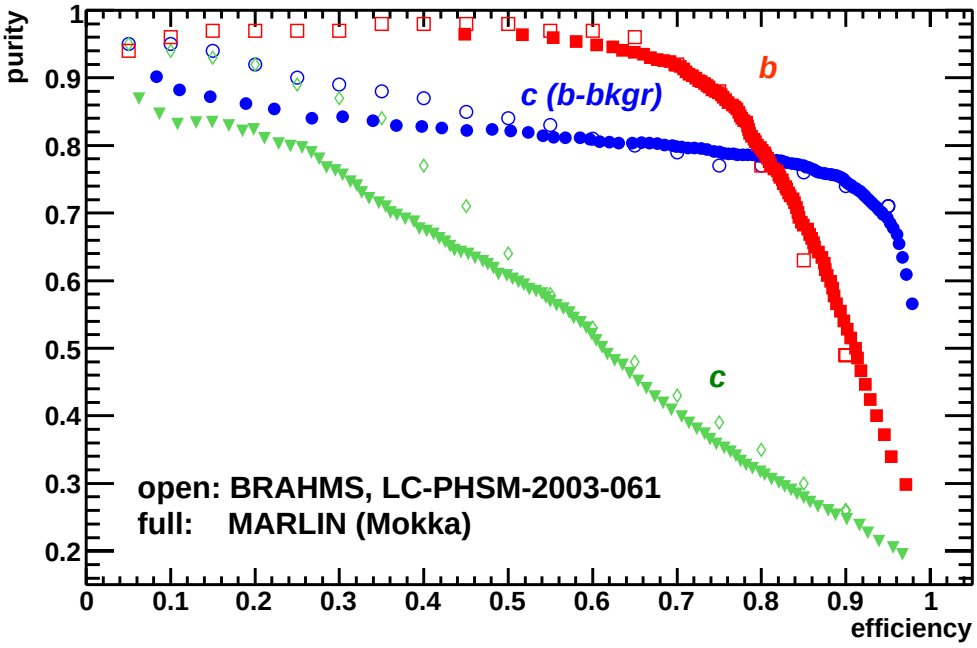

Figure 7 The purity of the flavour tags for b jets, c jets and c jets with b background only, plotted as a function of efficiency, for a centre-of-mass energy of 92 GeV. Full symbols correspond to the result from our package, embedded into Marlin and fed with input from Mokka, open symbols to the performance previously obtained with the Geant3-based MC simulation Brahms.

This system test was subsequently repeated with input events from the Mokka full MC, reconstructed using MarlinReco analysis routines. Initial discrepancies between results from the new code and former results from the Geant3-based full MC Brahms were traced to a problem with the track information input to the Vertex Package. Correction of this error produced results in agreement with expectations as shown in Figure 7. This is an example of how our code can provide useful feedback to the development of other ILC software packages and it is intended to continue the successful collaboration with other groups in this area.

For the ZVKIN branch of ZVTOP, the first release of the code provides the full algorithmic functionality, including access to the parameters of the algorithm, such as the minimum width of the ghost track. Since this branch of ZVTOP has so far only been used with the SLD detector, which had poorer track resolution in the main tracker, poorer impact parameter resolution in the vertex detector and restricted angular coverage compared to the ILC detectors studied by the four detector concept groups, it is likely that the parameters of the code will need to be tuned for the unexplored ILC environment. This tuning will be done in conjunction with the exploration of the use of ZVKIN for flavour tagging and quark charge determination.

The aim of the ZVKIN tests therefore was to convince ourselves that the algorithm was implemented correctly. This algorithm has two stages: finding the ghost track and using it to identify vertices, which may contain only one real track. The search for the ghost track relies on the fact that, due to the high momentum of the B hadron, the interaction point, the B hadron decay vertex and the subsequent D hadron decay vertex lie approximately on a straight line. The ghost track direction is initially set to the direction of the jet momentum and then optimised to approximate the B flight direction. Figure 8(a) shows that this optimisation is successful in decreasing the angular distance between ghost track and the B hadron direction. We also compared the decay lengths of the vertices found in the second part of the algorithm to the MC values, as shown in Figure 8(b). Given that parameters have not yet been tuned and that some hadrons decay very close to the IP and are hence difficult to find by any algorithm, fair agreement between reconstruction and MC values is seen for a reasonable fraction of the test sample.

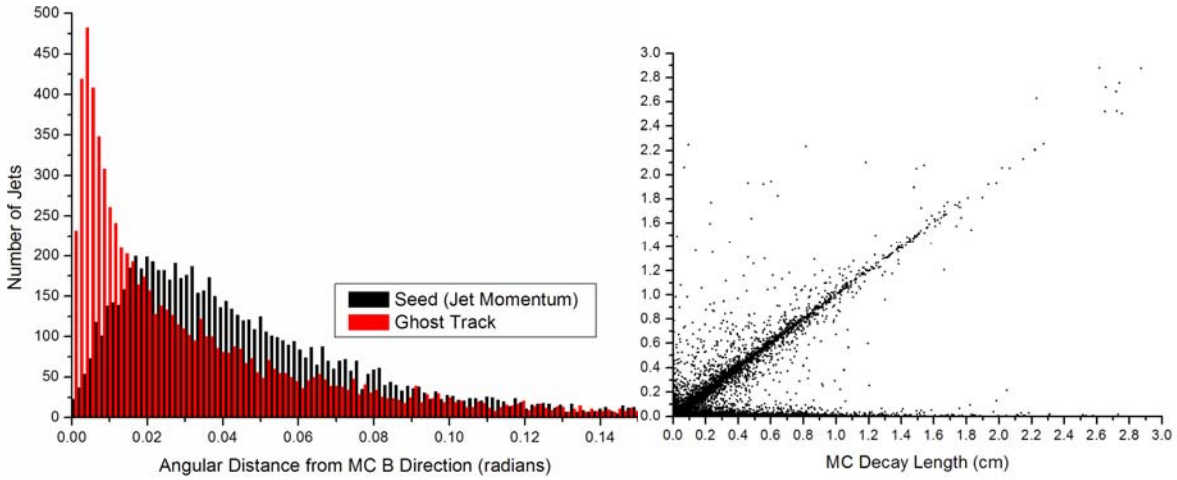

(a) (b)

Figure 8 Performance of the ZVKIN (ghost track) vertex finding algorithm: (a) shows the angular distance of the ghost track to the flight direction of the B hadron, compared to that between the B and the jet momentum; (b) compares the reconstructed decay length with the corresponding MC value.

Other physics studies

Quark charge selection and vertex detector parameters

The LCFI group has investigated the possibility of measuring the charge of hadrons containing heavy quarks at the ILC and hence performing quark charge selection. This can be done by assigning all the charged tracks in a jet to a primary, a secondary or perhaps also a tertiary vertex and then summing the charge of the tracks associated with the decay chain. Preliminary results of such a study are illustrated in Figure 9, which shows the vertex charge reconstructed for charged and neutral B hadrons in 100 GeV b jets. For charged B hadrons, there is a strong correlation between vertex charge and the quark charge, whereas the case of neutral B hadrons will require further analysis, including the "charge dipole" approach pioneered by the SLD Collaboration, and will always be limited by the effects of $B^0\overline{B}^0$ mixing. One measure of the efficacy of the quark charge selection is the leakage rate λ_0, i.e. the probability of reconstructing a neutral B hadron as charged.

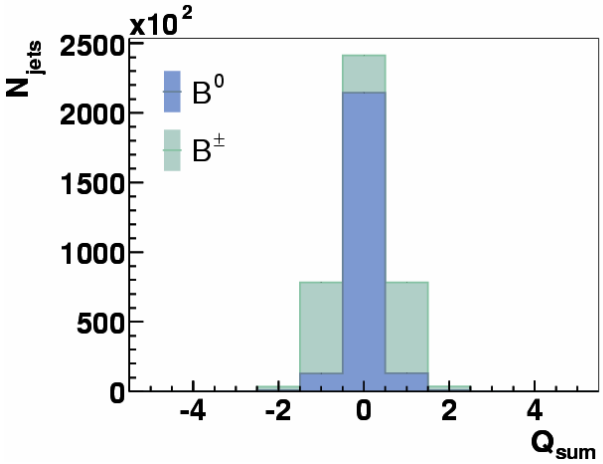

Figure 9 Preliminary vertex charge distributions for 100 GeV b-jets.

As quark charge selection requires the correct assignment of all tracks to secondary or tertiary vertices, the LCFI group suspected this may be particularly sensitive to changes in the vertex detector design, such as increases in the radius of the beam pipe. Hence, the effect of such changes was investigated. The results of this study are shown in Figure 10(a). The leakage rate, λ_0, is seen to increase as the beam pipe radius R_{bp} is increased from 15 to 25 mm or if the thickness of the vertex detector layers increases from 0.1% X_0 to 0.4% X_0 per layer. Increasing the minimum momentum of the tracks assigned to vertices by ZVTOP from 0.1 to 0.2 GeV also leads to a notable degradation in performance. The results emphasise the need to push all these parameters to their limits, since in the real detector all these effects will add up.

While from the change in leakage rate these effects may appear to be small, it was shown that the detector with increased beam pipe radius would require an increase in integrated luminosity by about a factor of 2 in order to compensate for the loss in charge identification performance for physics processes relying on independent measurement of vertex charge in two jets, as is illustrated in Figure 10(b).

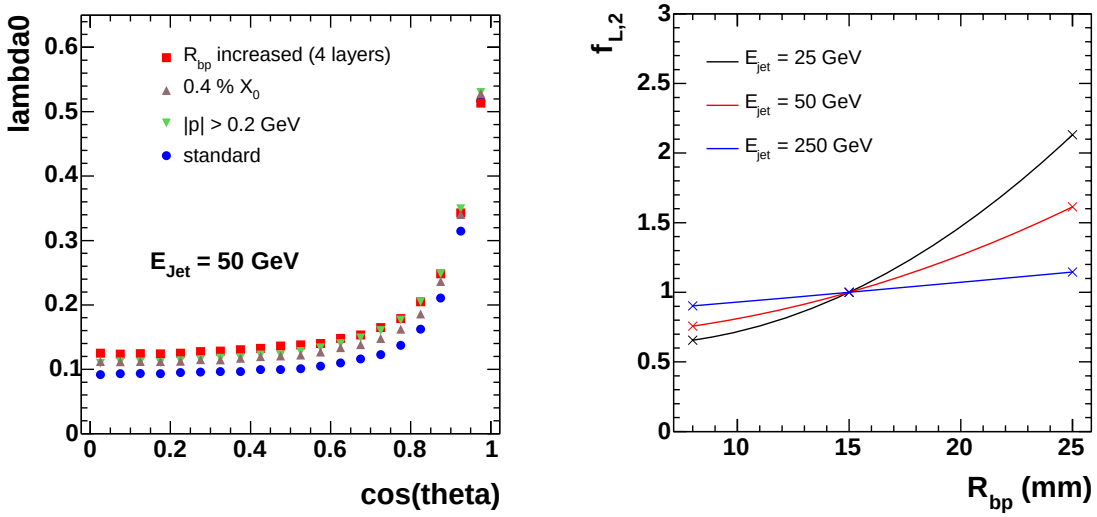

(a) (b)

Figure 10 (a) Leakage rate as a function of polar angle for various detector designs and reconstruction procedures. (b) Factor by which integrated luminosity has to vary to reach same physics sensitivity as the standard detector design for processes requiring vertex charge measurement for 2 jets, as a function of beam pipe radius.

Studies of effects of proposed changes to ILC design

Following the presentation of the ILC design at the Vancouver Linear Collider Work Shop in July 2006, several changes were proposed in an effort to reduce the cost of the machine without damaging its physics capability. One of these changes was termed the "low P" option. The suggestion was that initially only half the klystrons would be deployed in the ILC Linacs. This results in a halving of the luminosity, which can be recovered by further compression of the e^+ and e^- bunches at the IP. Studies of the resulting background at the IP in which LCFI was involved revealed that this option is incompatible with the current beam pipe and vertex detector designs. As is shown in Figure 11, the additional beam disruption increases the transverse momentum of the e^+e^- pairs produced when the bunches interact and causes these to get uncomfortably close to the beam pipe. Increasing the radius of this to provide the necessary safety margin means the inner radius of the vertex detector must be increased from about 16 mm to 24 mm. As demonstrated above, such an increase results in a significant loss of vertex detector performance which has a severe impact on the physics capability of the ILC. This argument led to the realisation that further compressing the ILC bunches is not desirable. The low P option is thus only being considered in the sense that, should insufficient klystrons be available at the start-up of the ILC, operation at reduced luminosity with nominal beam conditions should of be considered as a temporary measure.

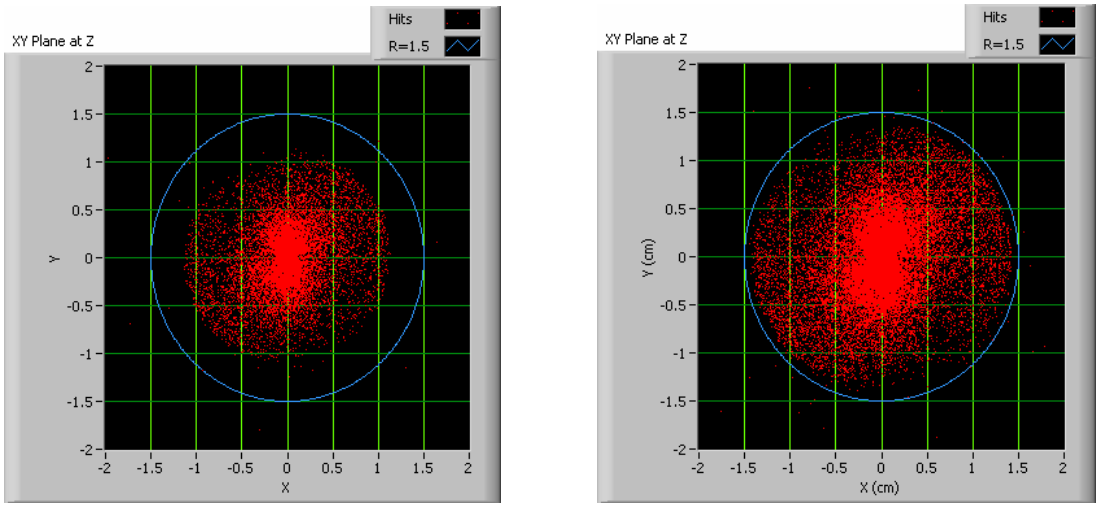

Figure 11 Positions of e^+ and e^- tracks produced in the ILC under nominal (left) and low P (right) operating conditions; in the latter case, the tracks are unacceptably close to the beampipe, indicated by the blue circle.

Future plans

Following the release of the Vertex Package, LCFI is in a position to take up the study of benchmark processes relevant to the vertex detector design. In parallel, we will continue to work on the tools in areas that directly feed into these studies.

With the recent additions to the MarlinReco reconstruction framework, this software environment has reached a level of maturity that will permit us to make the transition from the fast MC simulation and reconstruction SGV to the full MC program Mokka and a more realistic reconstruction.

A prerequisite for this transition is a more detailed study of the flavour tagging purities and of leakage rates for vertex charge reconstruction achievable with full reconstruction in MarlinReco as compared to SGV. In the testing of the Vertex Package, we used the "cheater" algorithms implemented in MarlinReco, in which tracks are fitted using a set of hits known from MC information, to check the performance of our code. This is similar to what is done in SGV. It will be necessary to study how the performance changes when this simplified approach is replaced by the full pattern recognition and track fitting software under development at MPI Munich[10]. These studies will permit us to give useful feedback to these developments.

Another obvious area for further study is the full exploration of the improvement that can be achieved by the use of the ghost track algorithm at the ILC. These improvements will affect both the flavour tag and vertex charge determination. The ZVKIN algorithm contributes to the charge dipole procedure, from which the quark charge can be obtained for short-lived B_d^0 decays.

As well as continuing tool development, we will take up the study of benchmark physics processes sensitive to the vertex detector design. We will revisit the study of Higgs branching ratios performed for TESLA[11]. The availability of earlier results, obtained with the GEANT3-based MC simulation and reconstruction software, Brahms, will permit comparisons of performance at the level of physics results, which

will provide a useful consistency check not only of the new Vertex Package but also of other parts of MarlinReco.

We will also work on a study of the process $e^+e^- \to b\bar{b}$, which will provide excellent sensitivity to new physics at the ILC, as shown e.g. by J. Hewett and S. Riemann[12]. In contrast to the studies by Hewett and Riemann, our study will take into account the experimental limitations that arise from realistic flavour tagging, as well as from quark charge sign selection. Pioneered for the ILC by LCFI, quark charge selection will for the first time be included in an ILC study in this work.

Both these processes are among the benchmark channels that the ILC community has agreed to focus on at the current stage of R&D[13]. We propose to investigate a third process in addition to the above and are currently discussing the various possibilities available.

WP2 – Sensor Design and Production

Introduction

The aim of the LCFI sensor studies is to design and produce a sensor that has performance adequate for the VXD of the ILC. We are developing two sensor types, the column parallel CCD (CPCCD) and the In-situ Image Storage Sensor (ISIS). The investigations of the Flexible Active Pixel Sensor proposed in the previous LCFI report to the PRC were not funded. The vertex detector of the ILC requires sensors of a length of about 10 cm and a width of about 1.5 cm (actual values depend on the precise location of the sensors). In order to avoid excessive occupancy in the inner layer of the VXD, with the associated track finding problems, the CPCCD must be read out up to twenty times during an ILC bunch train, i.e. at frequencies of up to 50 MHz. This is perhaps the major challenge LCFI faces in developing this sensor. Previous PRC reports have described the successful manufacture and testing of the first generation of CPCCDs, CPC1. Here we describe progress with the manufacture of the CPC2 and the design of test CPCCD structures (CPC-T) aimed at reducing the inter-gate capacitance and drive voltages of CPCCDs. Progress with the testing of the CPC2 is described in the WP5 section of this report.

The ISIS, in contrast to the CPCCD, stores signals in-pixel twenty times during the ILC bunch train and can then be read out relatively slowly in the long inter-train gap. The challenge here is to construct a storage CCD that is small enough to allow twenty cells to fit in the area of an imaging pixel, together with the necessary control circuitry. Here, we report on progress with the programme to design a second generation device, the ISIS2. Again, we describe the tests of the ISIS1, a "proof-of-principle" device constructed by e2v, in the WP5 section of the report.

CPCCD and ISIS design and manufacture

LCFI received 3 diced wafers of CPC2 and ISIS1 devices from e2v in mid-October 2005. The chips from one of these wafers are shown in Figure 12. Two further wafers were shipped to VTT for bump-bonding to our Column Parallel Readout (CPR) chips.

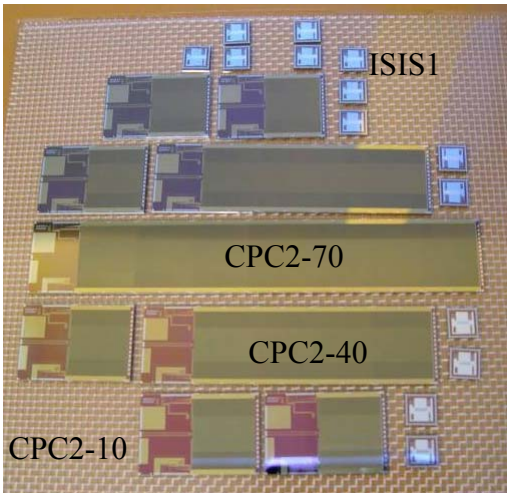

Figure 12 The sensors resulting from dicing one wafer; sensors from three such wafers were received by LCFI from e2v in October 2005.

The CPC2 chips are of sizes varying from an imaging area of 13 × 15 mm^2 (CPC2-10) to 92 × 15 mm^2 (CPC2-70).

Following the successful tests of the low speed version of CPC2-10, described later in this report, LCFI instructed e2v to continue with the manufacture of the high speed busline-free double level metal CPC2 devices. Two wafers were processed, diced and delivered to LCFI. A photograph of one of these diced double level metal wafers is shown in Figure 13. Further wafers will be finalised after feedback from the tests of the high speed bus line free CPC2 and of the bump-bonded assemblies of the CPC2 and its readout chip the CPR2.

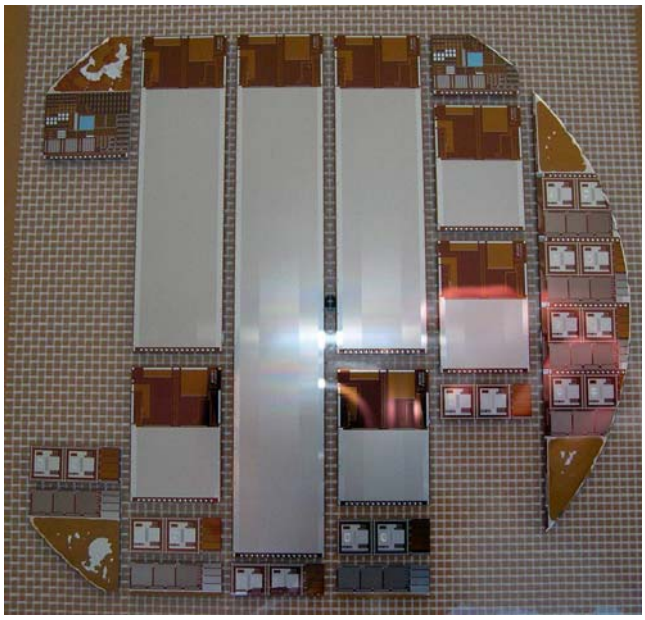

Figure 13 Double level metal CPC2 wafer.

Design of test CCDs

LCFI invested a lot of effort in the design of a number of small test CCDs (CPC-T) which will be used to investigate possibilities for reducing the power required to drive CPCCD readout. Ten different designs are now ready which will be produced with various process modifications, giving 29 different CPC-T chips in total. Two of the CPC-T designs are dedicated to investigating CPCCD operation at low clock amplitudes and achieving a low inter-gate barrier by varying the dopant implantation or the dielectric thickness. Six designs explore various possibilities for the reduction of the inter-gate capacitance. Two further devices have only clock bus lines to allow measurement of the capacitance due to these lines on the test devices. Several of these CCDs are novel LCFI designs. Their development involved over 5000 CPU-hours of 3D simulations on the RAL Synopsys computer farm.

The test chips share a common design and size. They have 10 columns and approximately 500 rows of 20 × 20 μm^2 pixels. Four columns in the centre of the array are equipped with 2-stage source followers, allowing operation at frequencies of up to 50 MHz with an external load. The phase capacitance is estimated to be 150 pF, which can easily be driven at the highest design frequency. Because of the low gate capacitance, the on-chip clock distribution buslines are not expected to be challenging.

Devices for reduced clock amplitude studies

The two test CCDs for clock amplitude studies have rectangular or profiled gates and will be 2-phase devices based on the CPC2 design, but the 4 polysilicon gates will have separate connections. This will allow us to study the fundamental limitations on the amplitudes needed to reliably transport the charge at high frequencies. External offset voltages applied between a pair of gates can be used to create an inter-gate barrier of arbitrary height. Some devices will have low or no p-type inter-gate implant specifically for these studies. The results of the tests of the double-level metal CPC2 sensors discussed in the WP5 section of this report give us reason to believe that low inter-gate implants will indeed have the desired effect and can be produced. Questions remain as to the tolerances achievable in producing these implants that will be addressed using these test devices.

Another process variation will involve making the p-type inter-gate implantation before the oxidation of the polysilicon gates. Simulations show that this creates a much lower potential pocket in the inter-gate gap and can therefore allow lower clock voltages to be achieved. In the third process modification, the inter-gate barrier will be made by using silicon nitride of different thicknesses, which has the advantage of eliminating the need for implantation of very small dopant levels, which are difficult to control accurately.

Devices for reduced inter-gate capacitance studies

Reducing the capacitance between the two phases of the CPCCD clock could offer numerous advantages for the clock driver system and for the overall power consumption of the sensors. LCFI has formulated several ideas which have the potential to reduce the gate capacitance by up to a factor of 4. Around 70% of the CCD capacitance is caused by the inter-gate component C_{ig} and our efforts have concentrated on reducing this. Simulations have shown that C_{ig} does not depend strongly on the actual shape of the gate overlap (obtained from scanning electron microscope images) or on the gate thickness. In fact, C_{ig} is given almost entirely by the geometrical capacitance between the gate plates. This important conclusion opens up new directions for capacitance reduction.

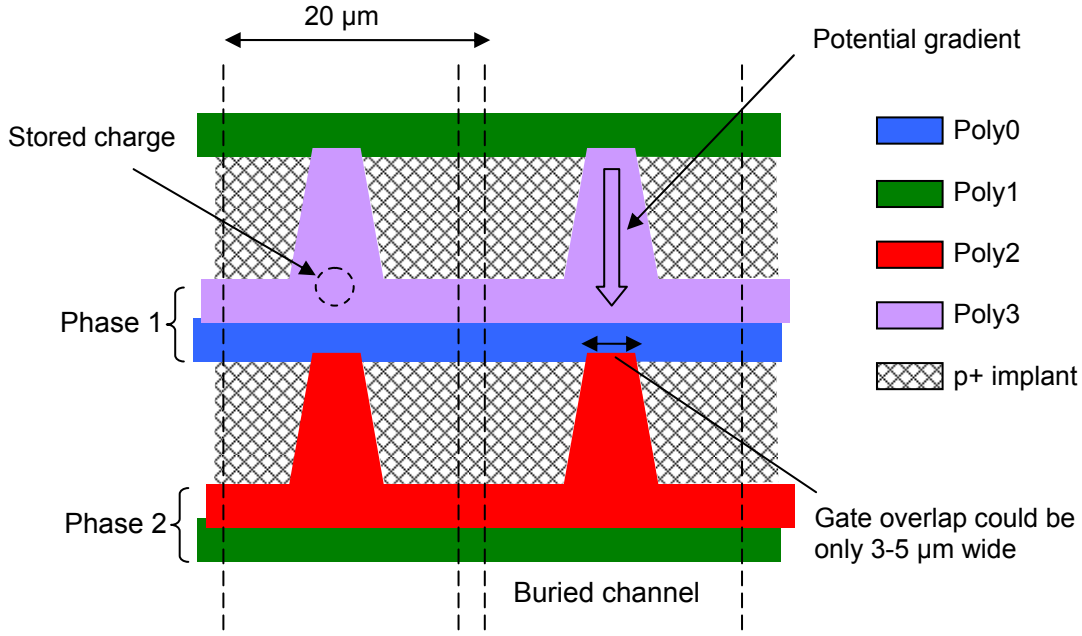

Figure 14 Schematic design of open phase CCD

One possibility is to keep the same shape of the gate overlap, but to reduce the gate area. Devices with such an architecture have been made by e2v for astronomy applications, where capacitance reduction is not pursued but comes as a by-product. Examination of existing capacitance measurements made on these open-phase devices (CCD22) shows that a factor of 2 reduction in C_{ig} is possible. The device uses profiled polysilicon gates to create a potential gradient in-pixel in the direction of the charge transfer. Shallow p+ implants are used to define the potential of the open silicon surface and to confine the charge to the buried channel under the profiled gates. Such a structure is illustrated in Figure 14.

The gate overlap between adjacent phases need only be 3...5 μm wide which would naively be expected to lead to a reduction of C_{ig} by a factor of about 4...7. In practice, the reduction is around 2, because the p+ areas are partially cut off from the substrate and serve as conductive planes between the two phases, adding to C_{ig}. The signal charges in the open-phase CCD are stored in a small volume at the base of the trapezoidal gates, which improves the radiation hardness with respect to bulk damage. This could be further strengthened by adding a narrow supplementary (or notch) channel. The open-phase CCD uses 3D potential effects to create the inter-gate barrier instead of implants, which inherently offers low clock amplitude drive.

Several further routes to low C_{ig} are under investigation. For example, the gate overlap over the channel stops of the standard 2-phase CCD can be removed and the buried channel made narrow to minimize the overlap. To reduce the capacitance to substrate, the oxide over the channel stops can be made thicker.

The scope of this work has expanded significantly over that originally foreseen, to the point at which a dedicated wafer run will be needed to explore all the options for clock amplitude and capacitance reduction. A much less demanding “passenger” CCD batch alongside the main production at e2v was envisaged in the previous LCFI report to the PRC. Nonetheless, the potential benefits of these studies are such that they are worth pursuing.

The production of the CPC-T test devices will start in spring 2007 and is expected to take 16 weeks.

ISIS developments

The ISIS, as is described in the last LCFI report to the PRC, is designed to store the signal charge in-pixel in a small CCD register, with readout taking place during the inter-train gap at the ILC. The ISIS1, a device deigned to provide a “proof-of-principle” of the ISIS concept was manufactured by e2v in two variants. One of these had a deep p well, the deep p implant being designed to shield the CCD register from the image charge in the pixel, the other did not. These designs were presented in our previous report. As is described later in this report, the ISIS1 chips without the deep p implant functioned well. However, it was not possible to turn on the output source followers of the sensors with a p-well as the doping levels were incorrect. e2v have agreed to manufacture the ISIS1 again with the correct p-well doping. Delivery is expected in April 2007. The new ISIS1 batch will consist of 9 wafers and will be split into 3 groups to minimise doping uncertainties. The doping levels were re-calculated by the LCFI team and at e2v. This was backed up by process simulations by an engineer at e2v.

One of the companies we are investigating as a potential partner in the production of the ISIS2 is DALSA. They manufacture small-pixel CCDs, and their process technology appears to be suitable for the manufacture of the ISIS storage cells. However, the DALSA CCDs are built with non-overlapping gates, which is typical of small-pixel CCDs and CMOS processes, and their surface radiation hardness has not previously been studied. A possible problem with these is that the trapped charge in the thick layers of SiO_2 above the exposed inter-gate gap could create potential pockets in the buried channel. This would be particularly visible at the low clock voltages used in the CPC and ISIS. Following discussions of this point with DALSA, the company supplied LCFI with five FTF2416 CCDs. These are 4 Mpixel image sensors with 9 μm square pixels.

The first step in the LCFI studies of the DALSA sensors was to verify that they can work at very low clock voltages. The supplied hardware was placed in a cryostat and modified to reduce the image clocks below the nominal operating level of 8 V. The CCD performed extremely well with clock voltages as low as 2.1 V, showing excellent noise performance of 20 e^- ENC and high quality X-ray response using a ^{55}Fe source, as shown in Figure 15. Charge smear due to insufficient clock amplitude became visible at 2 V.

Radiation damage effects were then evaluated. One chip was irradiated at RAL with a strong ^{90}Sr source, reaching a dose of 76 krad. No flat-band voltage shifts were observed in the gaps between the gates, and the device continued to operate with gate voltages of 2.1 V_{pp}. We do not consider further irradiation of this chip to be sensible due to the high dark current induced by the irradiation coupled with the difficulties of operating the DALSA electronics below –20 C. This makes measurements increasingly difficult. A further chip will be irradiated to cross-check the results from the first.

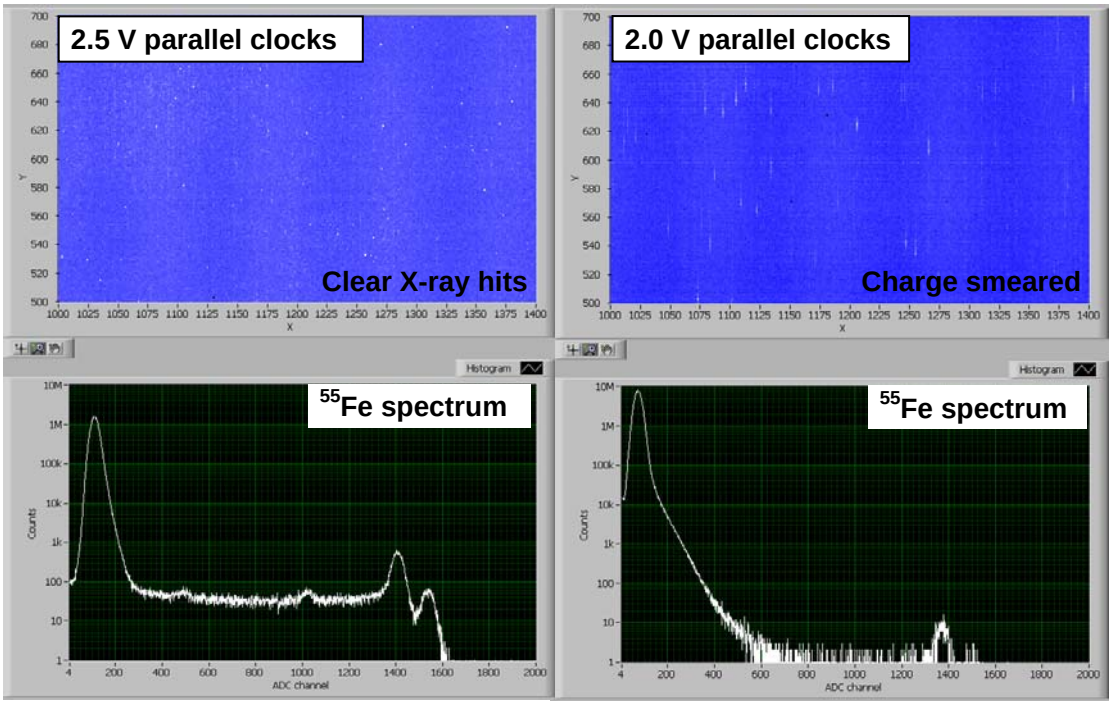

Figure 15 Performance of DALSA FTF2416 at very low parallel clock amplitudes. Even at 2.5 V, the Mn K_α and K_β lines are clearly resolved due to the very low readout noise. The serial clock frequency is 20 MHz.

The performance of the DALSA CCDs is encouraging and we are now discussing the possible design and manufacture of the ISIS2 with the company. A DALSA engineer has carried out a preliminary feasibility study for LCFI.

ZMD/ZFoundry, another vendor with promising capabilities, was also and contacted. ZFoundry has a 0.6 μm mixed CCD/CMOS process with two-level polysilicon and three-level metal. Buried channel CCDs with custom profiles and deep implants are also available. ZFoundry does not offer design services, unlike DALSA, so expertise in the Instrumentation Department as RAL has been sought to help with the CMOS aspects of the ISIS2 design. Work on this is now progressing.

Future plans

LCFI will first complete the evaluation of the bump-bonded CPC2 assemblies and the busline free CPC2. These studies will provide input into the design of the CPR2a chip. Once the manufacture of the CPR2a is complete, this will be tested and bump-bonded to CPC2 sensors. The bump-bonded assemblies will then be investigated.

The numerous CPC-T devices will be investigated. If it proves possible to reduce the inter-gate capacitance of the CPCCD using the architectures developed by LCFI, future large CPCCDs will exploit these architectures.

The chips produced in the re-run of the ISIS1 will be tested and the resulting information fed into the design of the ISIS2. Work towards this second generation ISIS will continue with discussions with the vendors DALSA, ZFoundry and also Jazz Semiconductor. The Collaboration aims to complete at least one ISIS2 design in the second half of 2007 and then to choose the most promising design for manufacture and testing.

WP3 – Readout and Drive Electronics

Introduction

The goal of this part of the LCFI programme is to design and produce the Column Parallel Drive (CPD) and Column Parallel Readout (CPR) chips necessary to clock and read out the CPCCDs. Clock signals of frequency 50 MHz (inner layer) and 25 MHz (outer layers) and of amplitude about 2 V_{pp} must be distributed across the entire imaging area of the CPCCD. As the capacitance of a device of the scale needed for the VXD is significant, currents of up to 20 A (outer layer CPCCDs) must be provided to achieve this. Progress with the design, manufacture and testing of the first CPD chip, the CPD1, is described here.

As previously reported, LCFI tests of CPR1 chips bump-bonded by VTT to CPC1 sensors have proven that our column parallel readout chips can successfully be bonded to CPCCDs with a column pitch of 20 μm and have demonstrated the feasibility of column parallel readout. LCFI has designed a further readout chip, the CPR2, which includes cluster finding logic and sparse readout circuitry, in addition to the amplifiers and 5-bit FADCs found on the CPR1. First tests of this chip are described in the following. Further improvements of the readout chip are planned and the design of the next generation chip, the CPR2a, is discussed below.

Column parallel readout

CPR2 studies

The design of the CPR2 chip is based on that of the CPR1 described in the previous PRC report, but major new features were introduced. The CPR2 includes the capability to apply on-chip thresholds and performs clustering and data sparsification. Following the CPR1 tests, several small changes were also made to the design of the CPR2. For example, the gain of the charge amplifiers was increased and a filter was introduced that sets a well defined frequency cut-off for the amplifiers. Testability was greatly enhanced by providing a test register and other structures. This makes possible the study of the analogue and digital performance of the CPR2 at different clock speeds. Variations from channel to channel across the full width of the chip can also be investigated much more satisfactorily than was the case with CPR1.

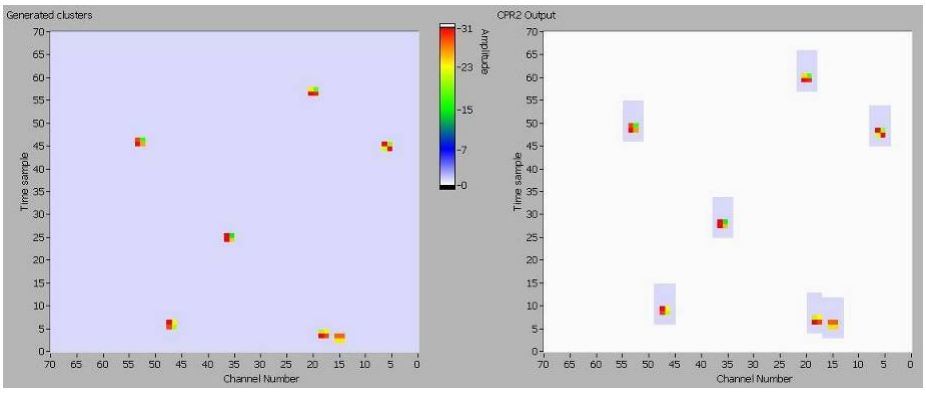

Figure 16 Test patterns input to the CPR2 (left) and the resulting output (right).

The first tests of the CPR2 were dedicated to the performance of the digital section of the chip, i.e. the cluster finding and data sparsification. Test patterns were input to the chip and then read out, as is illustrated in Figure 16. The results in the Figure show a case in which the chip functions perfectly. However, the tests revealed some dead-time problems. If successive hits are closely spaced in a column, the later hits are then not identified by the logic. This problem is understood and is addressed in the subsequent version of the chip, CPR2a. It will not prevent use of the chip for testing the CPC2.

Studies were also made of the performance of the analogue section of the CPR2. These again demonstrated that the chip functioned largely as hoped. Some minor difficulties were identified in the digitisation of test pulses, which demonstrate small hysteresis effects, i.e. the signal resulting from a given voltage input depends on whether the previous voltage was higher or lower than that being digitised.

The results of the laboratory tests of the CPR2 have been used to improve the simulations of the chip and hence a deeper understanding of the performance of the device has been obtained. For example, it is now understood that the abovementioned hysteresis effect arises from parasitic charge storage in elements of the ADC and that this can be combated in several ways in the CPR2a.

CPR2a design

The CPR2a incorporates changes in the digital logic for cluster finding and readout designed to combat the dead-time observed in the CPR2. This arises as the CPR2 suffers from a limited front-end memory depth. After a cluster has been detected in a column, it is necessary to read out that column before the memory can be freed up to store a new cluster. The CPR2a design has an extended column buffer which can store information on 3 clusters, including the time-stamp data. A data header register has been added in order to keep track of multiple clusters within the column.

This new architecture, illustrated in Figure 17, overcomes the dead-time problem for clusters which are closely separated in time. The new read and write processors provide flexibility in the formatting of data, for example avoiding unnecessary repetition of time-stamp data.

Development of the column parallel drive chip

The CPCCD Clock Driver ASIC, CPD1, is designed to provide the high currents necessary to charge and discharge the large CPCCD gate capacitances at frequencies of up to 50MHz. The challenges for the CPD1 are:

1. The CPCCD represents a capacitive load of up to 127 nF to ground per phase, with a 0.1 Ω distributed gate resistance.

2. The chip must provide a clock signal of up to 3.3 V peak-to-peak at 25 MHz.

3. In order to charge the 127 nF capacitance to a voltage of 1.65 V in 10 ns, a current of about 21 A per output is required.

4. The bunch train structure of the ILC requires that the CPD operate for 1 ms, with a duty cycle of 0.5%.

5. Operation at 50 MHz, with a reduced capacitive load of 40 nF, is required for the CPCCDs in the inner layer of the VXD.

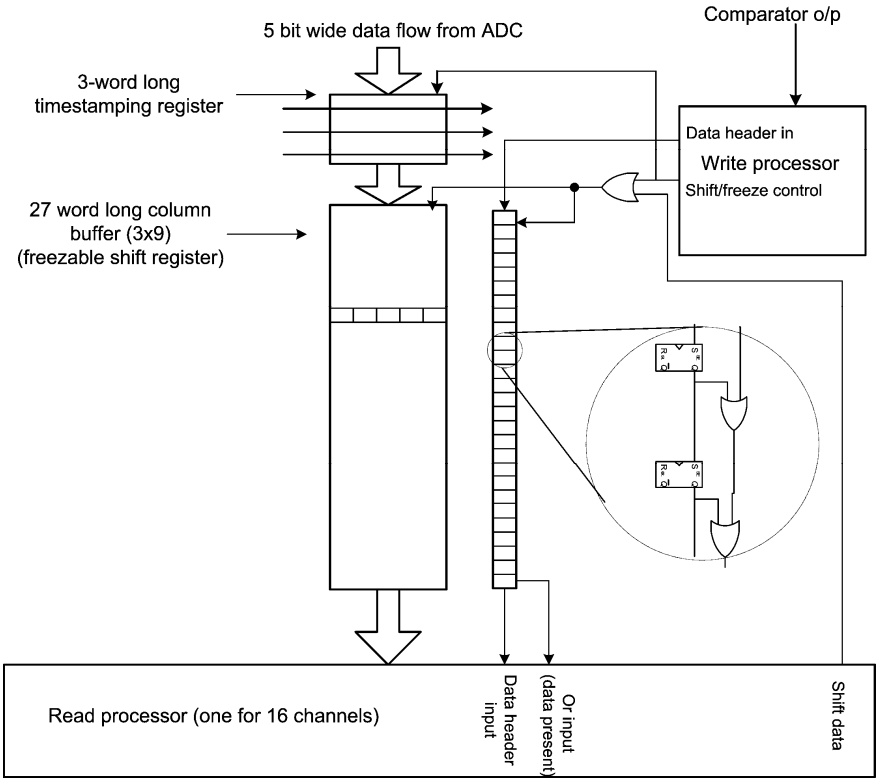

Figure 17 CPR2a block diagram.

The CPD1 is designed to have three modes of operation, suitable for different clock frequencies.

a) Inverter Mode, in which the gates are connected directly to the supply voltage and to ground via switch transistors. This results in the largest swing of gate voltage, with short time constants. However, there is no control over the output slew rate. This mode is best suited to 50 MHz CCD clocking.

b) Fast Mode, in which the gates are driven by source followers with variable input voltages. This gives access to frequencies from ~ 5 to ~ 25 MHz and allows a ramp wave form to be produced.

c) Slow Mode, in which again source followers are used and in which frequencies from ~ 1 MHz to ~ 10 MHz are generated with accurate control of the voltage slew rates.

The CPD1 is divided into 8 blocks, each with 4 output bond pads for phase 1 and 4 pads for phase 2. The blocks can be separately enabled, allowing the testing and operation of the driver with a wide range of load capacitors.

As external connections to the CPD1 can introduce significant impedance – the inductance of 1 mm of bond wire is about 1 nH – two of the outputs of the CPD1 are permanently connected to an on-chip capacitor of 2 nF. The connection is made through wide metal tracks, which gives much lower inductance than is achievable

with an off-chip capacitor. This allows testing of the driver performance for these channels without the limitations introduced by external bond wire inductance.

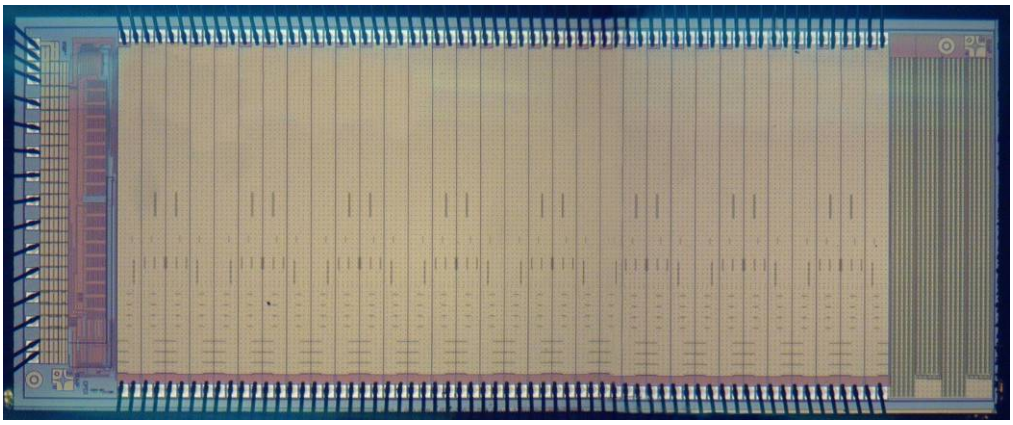

Figure 18 Photograph of a CPD1.

The driver chip uses two digital-to-analogue converters (DACs) to provide references for slew rate control – one for the positive edges and one for the negative edges of the clock pulses. The DAC input bits are stored in a serial register which also controls the selection of Inverter, Fast or Slow Modes and enables or disables the various blocks of the chip. This register is loaded in parallel from a second register which in turn is loaded sequentially from the external data input line. The parallel load function ensures that all of the internal control bits change simultaneously, so the driver is always in a well-defined state. The driver has a power-down feature, which reduces the static current to zero when the clocks are inactive.

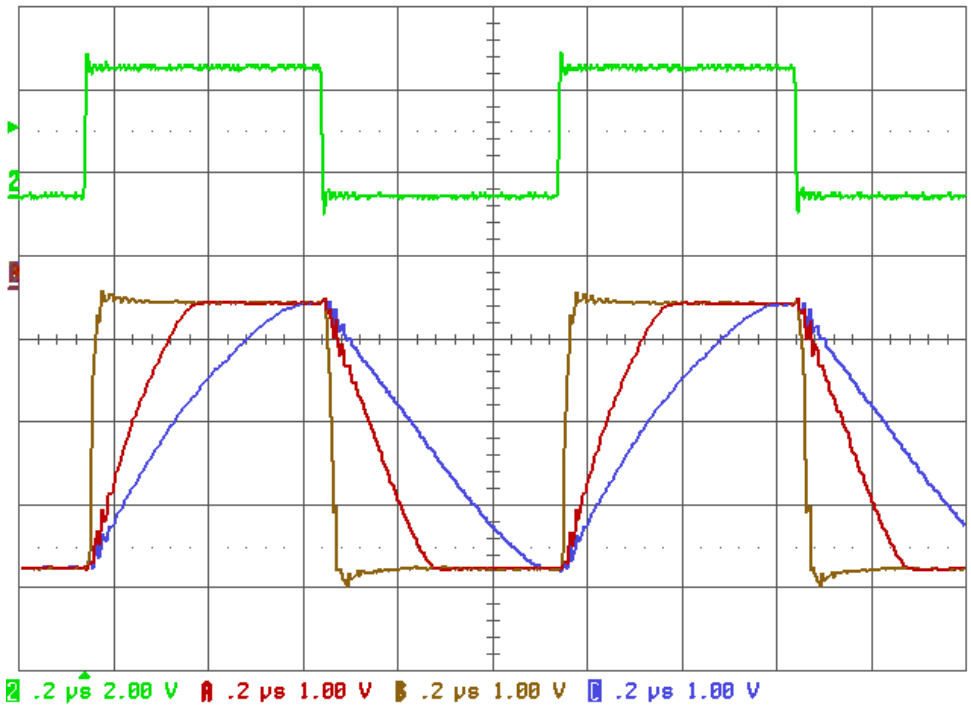

Figure 19 Slew rate control for Slow Mode.

CPD1 manufacture and testing

The CPD1 design was submitted for manufacture in an AMS multi-project wafer run. Forty chips were delivered to LCFI, one of which is shown in the photograph in Figure 18. The three distinct regions of the CPD1 layout are visible in the photograph. The control circuitry is on the far left; in the centre is the array of drive transistors, with bias circuitry and local decoupling capacitance; on the right is the 2 nF test capacitor hardwired to one pair of drive transistors.

Three CPD1s were bonded into Pin Grid Array packages for initial testing on the IMS chip tester at RAL. Although these packages provide full connectivity of the multiple power supply and output pads, the bond-wire lengths are excessive (6...7 mm) and there are also large parasitic inductances in the package itself. This format is thus not ideal for high frequency testing of the full driver with an external capacitor. However, low frequency tests are possible, particularly with the on-chip 2 nF capacitor where output inductances are small.

The preliminary test results using this set-up demonstrate that the CPD1 performs well over a wide range of load capacitances and clock frequencies. The driver waveforms can be controlled accurately over a range of slew rates, as is shown in Figure 19. In this Figure, the top trace is the Phase 1 Clock Input, with the driver output shown below for three DAC settings of the Slow Mode of the chip. The rise-times have been successfully adjusted over the range 20 ns to 400 ns.

Further tests were performed with the CPD1 wire-bonded into the dedicated test board described in the WP4 section of this report and shown in Figure 22. In addition to tests with the on-chip capacitance, this allows connection to a low inductance capacitive load via a smaller inductance than is present when using the Pin Grid Array packaging. First results of these tests are shown in Figure 20, in which the voltages produced across an external load of 40 nF at a frequency of 25 MHz are illustrated. As can be seen, even with significantly larger inductances than will be present when the CPD1 is bump-bonded to a sensor, voltages adequate for driving the CPCCD are achieved.

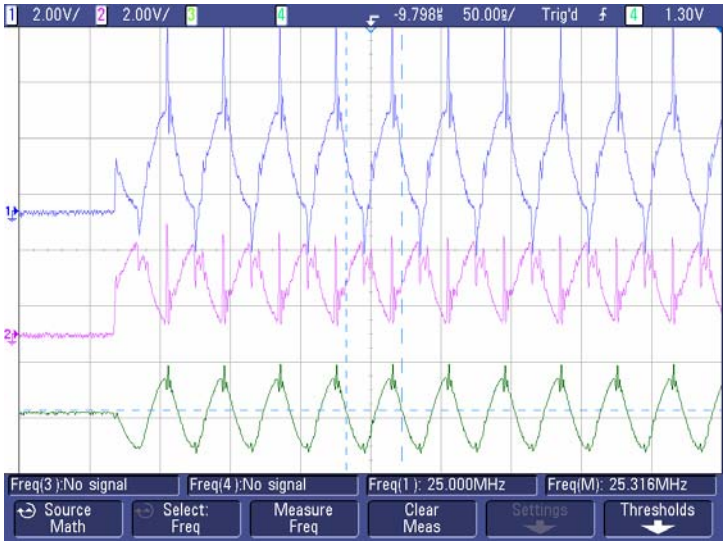

Figure 20 Drive signal with 40 nF load for phase 1 (upper blue trace), phase 2 (central purple trace) and difference of the two phases (lower green trace).

Future plans

Immediate future plans include the continued and extensive testing of the CPD1 and the completion of the design of the CPR2a. This latter will take account of the results of tests of the CPR2 and will also involve the further development of visualisation and simulation tools, so that simulated cluster data, incorporating latest knowledge in the expected backgrounds at the ILC, can be injected into a full simulation of the readout chip and the results presented in graphical form, allowing the influence of changes in the chip design on the expected chip performance to be easily evaluated.

WP4 – External Electronics

Introduction

The testing of the ISIS and CPCCD sensors requires electronic circuitry which provides all the control, biasing, readout and monitoring necessary for the operation of both the sensors and the readout and drive chips. These functions are provided by various motherboards and their associated electronics, the design and production of which is the subject of this work package and of this section the report. Boards for the independent testing of the column parallel drive (CPD) chips have also been developed and constructed. The design of motherboards for the investigation of the CPC-T low voltage and capacitance test structures is nearing completion.

Design and production of CPCCD motherboards

The efficient propagation of clock signals into and across the CPCCD becomes increasingly challenging as both the frequency and the sensor size increase, due to the large capacitance of the sensor. Large currents are required to obtain the voltages needed across this capacitance and even small resistive and inductive impedances in the path of that current can have severe consequences. Minimization of the inductance of the clock connections is an important goal which can be achieved, for example, by exploiting the cancellation of the magnetic fields induced by the opposite currents of the two clock phases.

Several techniques for producing the clock signals are being investigated. It is hoped that the custom CPD chip, described in the WP3 section of this report, will ultimately provide the drive for the full scale and full speed CPCCDs, but the use of miniature air-cored planar transformers, implemented in a multi-layer PCB, and a commercial driver chip is also being investigated for test purposes.

The 10-layer CPCCD motherboards provide all the bias voltages for the CPCCD, high current clock signals, analogue readout for 12 CCD columns, the digital interface for two types of readout chips, temperature monitoring and also a recess in which the CPCCD sits surrounded by fine-pitch bonding pads. The motherboard MB 4.3 uses CCD clock drivers based on transformers designed to drive the double-metal CPC2 chips. MB 4.4 uses MOSFET drivers, and MB 5.0 is the first board for which the clock signals will be provided by the Collaborations drive chip, the CPD1.

Two transformer-based MB4.3 boards were produced for the testing of the high speed double-level metal CPC2 sensors. As is described in the WP5 section of this report, these sensors were successfully read out at 45 MHz. However, the tests showed that the two phases of the clock have somewhat different amplitudes and are not always in perfect anti-phase. As is shown in Figure 21, calculations made using Spice have shown that this behaviour can be accounted for by parasitic capacitance between the primary and secondary windings of the transformer. Improvements to this drive system are currently being designed.

Two MB4.4 boards have been assembled. One of these is used for analogue measurements with the single-level metal CPC2. The second is used in tests of a bump-bonded CPC2/CPR2 assembly, as is described in the following section of this report.

The design of MB5.0, which will be the first motherboard to use the CPD1 driver chip, is complete. This board uses a new grounding scheme which significantly reduces the number of components and replaces the rather cumbersome cabling system that has been used to date with four 50-pin flexible flat cable (FFC) connectors. All cabling is thus uniform, low mass and can be easily disconnected.

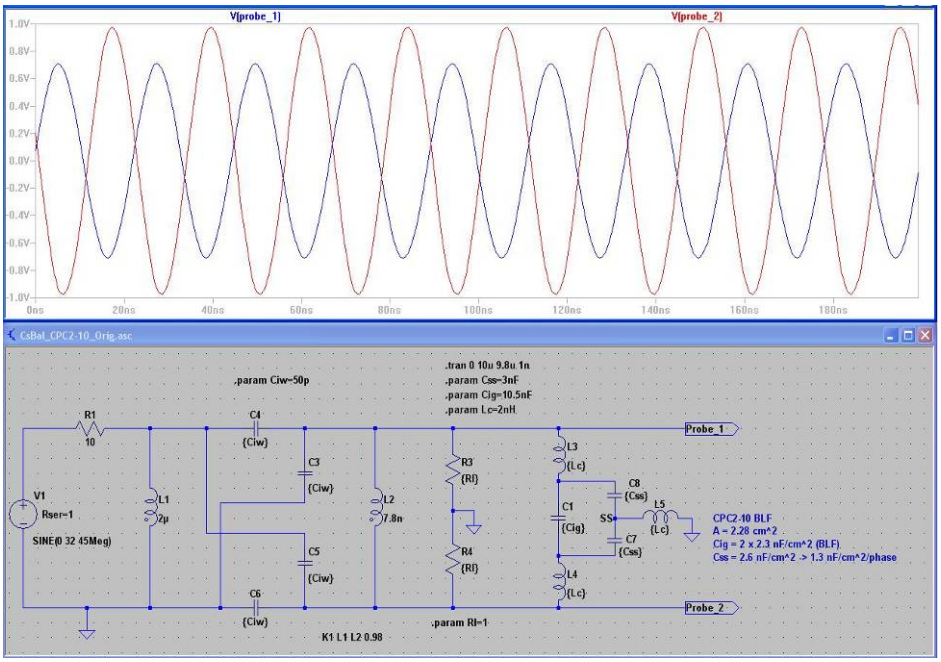

Figure 21 Spice model of the transformer drive circuit and the clock signals produced by the circuit.

VME module design and construction

A further essential component of the test electronics is the BVM2. This provides a standard, VME based, means of communication and control for all the test boards. The BVM2 daughter cards take care of different specific applications (such as providing programmable delays etc.). Ten BVM2 modules and several daughter cards have been produced and distributed for use at the LCFI test stands.

CPD1 test board

The test board for the CPD1 driver chip has been designed and produced. The test board allows the connection of a capacitive load with low inductance to the CPD1. The board is controlled by the BVM2 VME module described above, which downloads control information to a shift register in the CPD1 through a serial link and also provides all the necessary timing signals.

The first CPD1 has been mounted onto a test board and tests of the chip have started. Figure 22 shows part of the test board, with the CPD1 connected to an array of low inductance capacitors. These capacitors can be replaced by a custom capacitive load with low inductance which is being constructed in the thin film facility in Oxford.

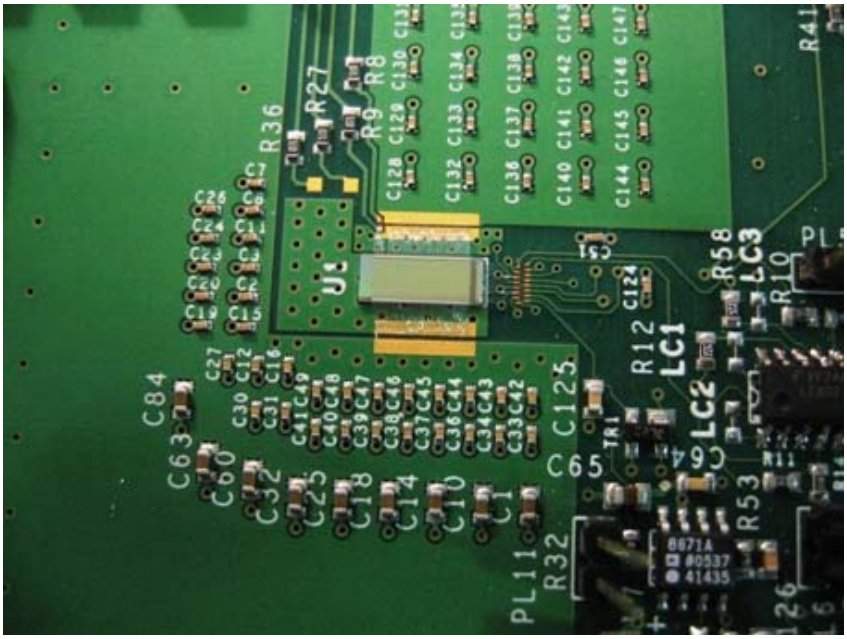

Figure 22 CPD1 driver chip mounted in its test board.

Future plans

This Work Package will continue to support the development and testing of the CPCCD with its readout and drive chips. The next set of motherboards to be produced will be those for the study of the CPC-T low capacitance and low drive voltage test structures. Motherboards will also be designed that will allow operation of CPCCDs in the EUDET test beam.

WP5 – Integration and Testing

Introduction

Since the last report to the PRC, LCFI has received both single and double-level metal versions of the CPC2 from e2v, the latter of which are designed for high speed operation. Here, the results of the tests of these devices are described. The Collaboration has also received and tested its first ISIS sensor, the ISIS1. Progress has also been made with the production of hybrid sensor/readout assemblies. CPC2 sensors have been successfully bump-bonded to CPR2 chips by VTT. First results from the test of a CPC2/CPR2 assembly are presented in this section of the report. Radiation damage studies have continued with the creation of a Synopsys-TCAD model of the CPC2 and the establishment of a test facility for investigating irradiated CPCCDs.

CPC2 tests

Following the receipt of a batch of CPC2 sensors from e2v, tests were first performed of two single-level metal CPC2-10 devices, one made on a 100 Ω.cm, 25 μm thick epitaxial layer and another on a 1.5 kΩ.cm, 50 μm layer. These were single-level metal sensors, not optimised to work at high speed but representative of the design and the process at e2v. Figure 23 shows the performance of one of these CPC2-10s at a 1 MHz column parallel drive frequency and with 5 V square-wave clock signals. These results were obtained using a commercial MOS gate driver chip, which is able to deliver 8 A peak currents to each phase. The maximum clock frequency achievable with this configuration is about 10 MHz for the single-level metal devices and about 5 MHz for the double-metal variants.

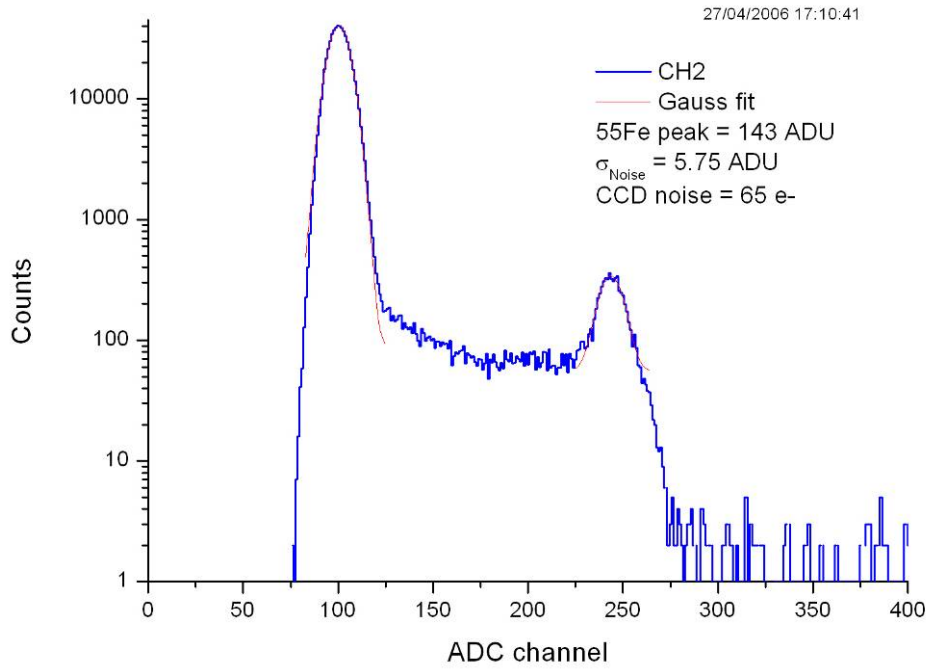

Figure 23 The first ^{55}Fe spectrum obtained from one output of CPC2-10 at 1 MHz clock frequency.

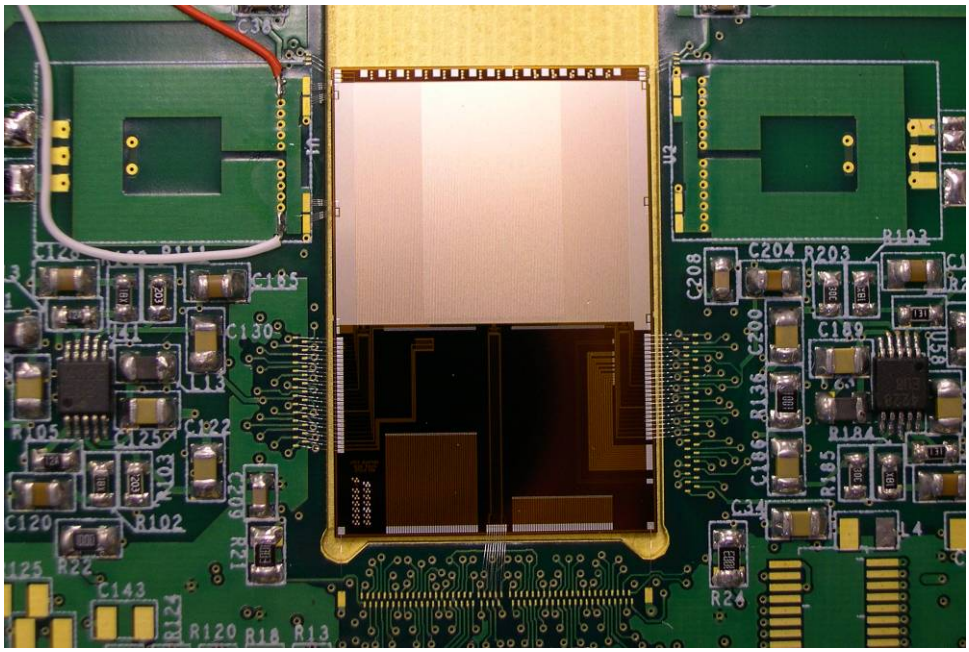

Figure 24 Double level metal CPC2-10 in MB4.3.

First tests have also been made of the busline free (double-level metal) CPC2-10 devices from this batch of sensors. The sensors were mounted in a MB4.3, as is illustrated in Figure 24, and driven via the transformers on the motherboard using a power RF amplifier.

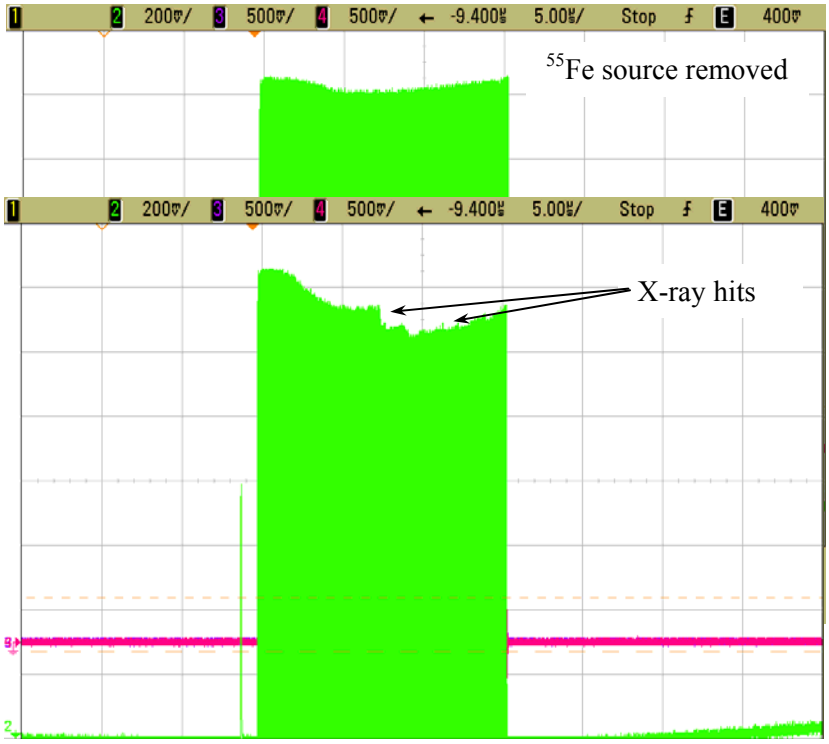

Figure 25 Oscilloscope traces of a discrete 2-stage CPC2-10 output at 45 MHz clock frequency with and without an ^{55}Fe source.

X-ray signals from an ^{55}Fe source were observed at clock frequencies of up to 45 MHz, as is shown in Figure 25. The clock amplitude used was ~ 2 V_{pp}.

During these tests, it was observed that the phase relationship between the two clock signals was not perfect and that the clock amplitudes were not always well matched. This was found to be due to the effects of parasitic capacitance in the transformers, as is described in the WP4 section of this report.

A surprising discovery during the tests of the double-level CPC-10 was that efficient charge transfer is still observed at clock amplitudes as low as 0.4 V_{pp}. Subsequent investigations revealed that this is due to a very low inter-gate implant barrier having been used in the production of these sensors.

Construction and testing of CPC2/CPR2 assemblies

Following delays due to a fire at their production facility, VTT delivered 18 CPC2 sensors bump-bonded to CPR2 chips to LCFI. The sensors used in this shipment are single-level metal devices. VTT still have a number of CPR1 chips which can be used for further tests if needed. Six bump-bonded CPC2-10/CPR2 chips were DC tested in order to eliminate faulty assemblies. No problems of the type previously observed with CPC1/CPR1, i.e. shorted substrates or CPR power supplies, were discovered. Several assemblies were wire bonded to motherboards, one of which is illustrated in Figure 26.

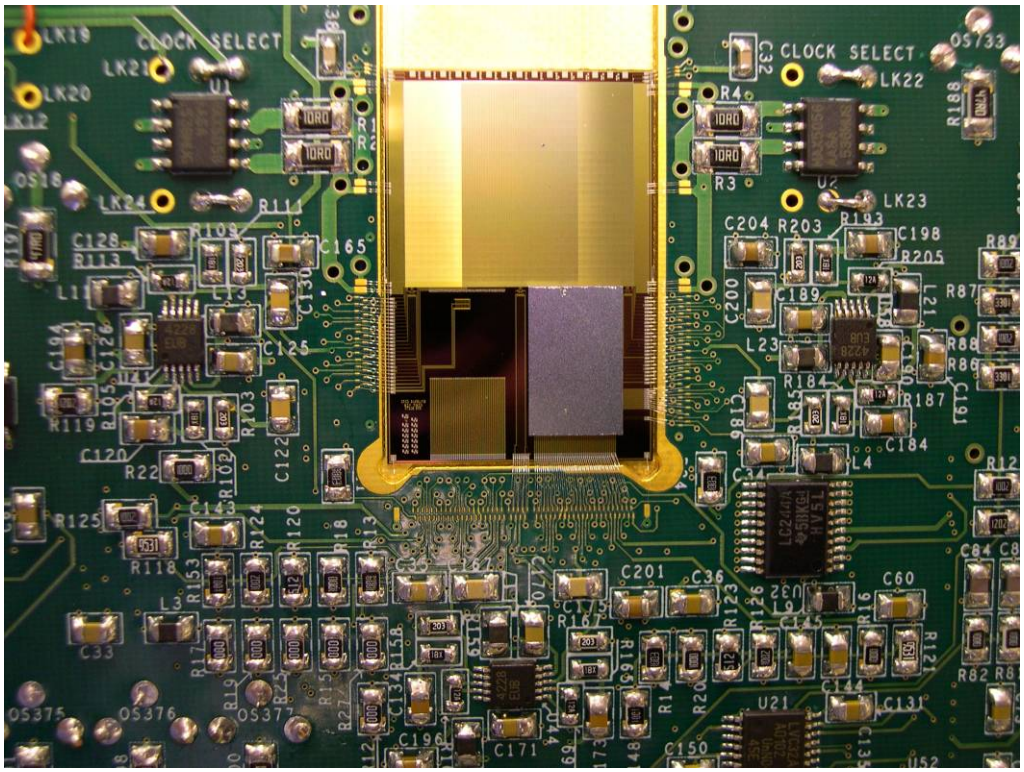

Figure 26 CPC2-10 bump-bonded to CPR2 in MB4.2.

The results of the first tests of this bump-bonded assembly using an ^{55} Fe source are shown in Figure 27. The X-ray signal is clearly observable, demonstrating the functioning of the CPC2/CPR2 hybrid.

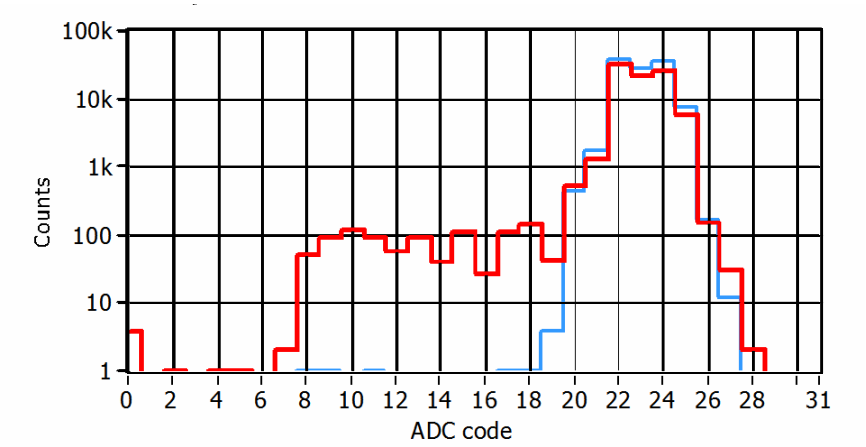

Figure 27 Data from a bump-bonded CPC2/CPR2 assembly with (red line) and without (blue line) an ^{55}Fe source.

Capacitance measurements

Accurate measurements of the phase capacitance of the e2V CCDs are important in order to tune the simulation tools used for the study of capacitance and its reduction in WP2. Measurements on CPC2 were carried using two different techniques.

The first method relies on measuring the time constants for the charge and discharge of the gate capacitance, clocked by square waves from a driver through a resistor. This measurement is done using a working CCD at low temperature, thus eliminating possible errors due to charge accumulation in the potential wells. The method measures the total gate capacitance, consisting of the inter-gate capacitance C_{ig} and the gate-to-substrate component C_s. The latter is easily calculable from the device parameters and is subtracted to get C_{ig}.

The second technique is more versatile because it uses a precision LCR meter working at many frequency points and able to supply different DC biases. It is somewhat more difficult to get right, because the CCD may have some stored charge which distorts the measurement. It has the power to determine separately C_{ig} and C_s and does not rely on indirect measurements.

Both techniques have produced comparable results. These measurements will continue to be used for the study of the dedicated capacitance overlap structures on the CPC2 wafers. These are polysilicon gates designed with an overlap varying from +3 μm to –2 μm (i.e. a gap). Due to the effects of the various processing steps, the real overlap is not a linear function of the design overlap, especially around a value of zero. These structures will provide us with calibration data for the simulation work in WP2.

Radiation hardness studies

Radiation hardness calculations using the Synopsis TCAD program previously made for the CCD58 and described in our earlier reports to the PRC have been extended to the CPC structure. The preliminary results of these calculations indicate that the temperatures at which the CPC is least influenced by the bulk damage caused in the CCD by the radiation doses expected at the ILC lie around 230 to 250 K, as shown in Figure 28.

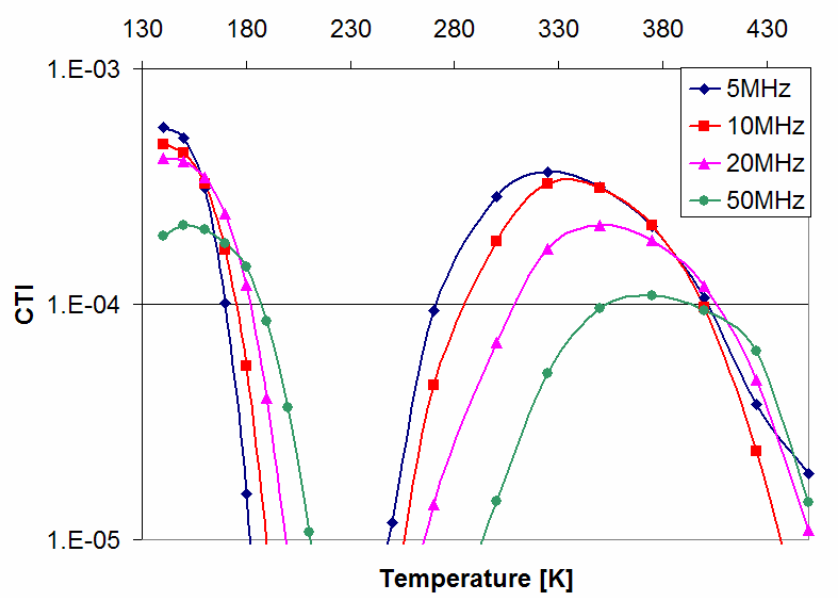

Figure 28 Charge Transfer Inefficiency (CTI) for the CPC as determined using the Synopsis TCAD program.

Test of the ISIS1

Study of the ISIS1 started immediately after delivery of the chips. First tests were made using a light source. These demonstrated that the ISIS1 does indeed see signals and that these are successfully stored in the in-pixel CCD registers of the device. Further tests of the ISIS1 were then performed using an Fe^{55} source. Typical results are shown in Figure 29, in which signals recorded in the five time slices of the ISIS1 and the resulting X-ray spectrum are shown. These results prove the feasibility of the ISIS concept for the detection of minimum ionising particles, which would cause the deposition of a similar charge.

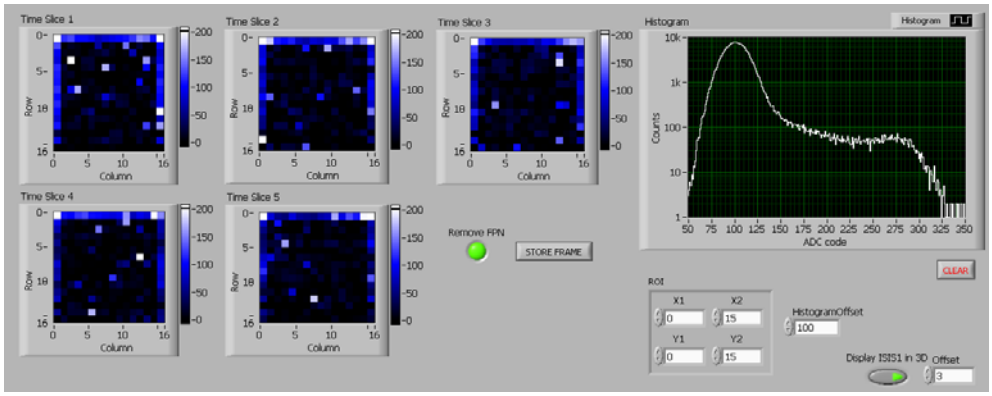

Figure 29 Fe^{55} signals from the ISIS1 showing signals in the 5 time slices recorded within each pixel. Note that the top row of pixels and the two side columns are not protected and collect diffusing charge, while the bottom row is protected by the output circuitry

Future plans

The busline free CPC2 will be evaluated, firstly using the present MB4.3 with its transformer-based drive systems then subsequently using MB5.0 and the CPD1. The

tests of the bump bonded CPC2/CPR2 are expected to be completed by mid-2007. The results from these tests will be used to tune the design of the CPR2a chip. Further sensors will be sent to VTT for bump-bonding after completion of the evaluation of the first batch.

The tests of the new ISIS1 chips with p-well will start as soon as these are delivered and the results will be fed back into the design of the ISIS2.

Tests of the numerous CPC-T devices will allow LCFI to determine if the concepts we have developed for reducing the inter-gate capacitance of the CPCCD can be realised in silicon. If these tests are successful, at least some future CPCCD designs will incorporate the new architectures so they can be tested in a full scale sensor.

WP6 – Mechanical Studies

Introduction

The aim of the LCFI Collaboration's mechanical studies is to develop the techniques needed to support the extremely thin silicon sensors that will be used in the VXD with a minimum of material, while providing the very high degree of mechanical stability necessary to allow precision measurements of particle tracks. Current investigations are concentrated on identifying the materials best suited to constructing sensor ladders and developing the methods necessary to assemble these ladders. Previously reported studies have shown that neither unsupported 50 μm thick silicon sensors, held under tension, nor 20 μm thick silicon sensors on a beryllium substrate can provide the required stability while keeping the material budget below the target of 0.1% X_0. The more recently studied foams have proved more promising, as is described below.

A further important issue for the VXD is the maintenance with minimal fluctuations of the operating temperature needed for the sensors. We must understand whether gas cooling, with gas flows at rates that will not induce mechanical oscillations in the sensors, is adequate and calculate the likely temperature fluctuations so we can determine the mechanical stability of the system. We are performing Computational Fluid Dynamics calculations of a test ¼ VXD model incorporating heating elements to simulate the sensors and their readout, and comparing the calculations with measurements made using the model. The goal here is to develop the tools that will allow us to reliably calculate the expected thermal characteristics of a variety of VXD structures and cooling schemes, allowing us to choose the best design with a minimum of prototyping.

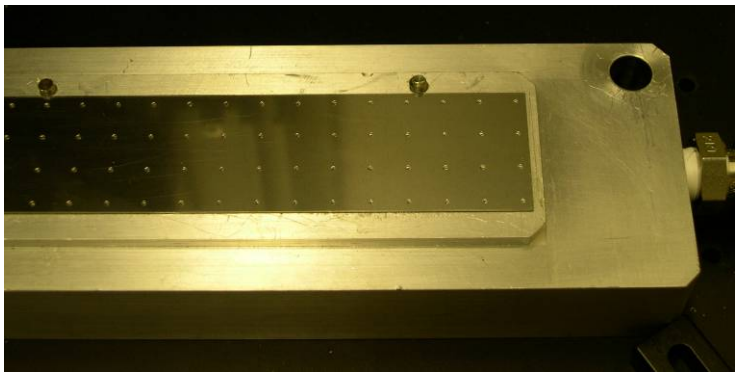

Figure 30 Stainless steel and glass ladder constructed using glue robot shown on vacuum chuck.

Ladder prototyping

LCFI has successfully commissioning two major new pieces of equipment, an EFD TT 525 glue robot and the Keyence LK-3010M laser displacement meter. The robot was programmed and used firstly to build a sandwich structure made from stainless steel and thin glass. It performed flawlessly, and the structure is shown (still on the vacuum fixture) in Figure 30. The precision glue spots joining the upper glass layer to the steel core are clearly visible. This build exercise also served as a test of ladder

assembly procedures, and was followed by the construction of a reticulated vitreous carbon (RVC) foam core ladder which is currently being surveyed.

The new laser has been integrated into the cryogenic surveying system and is performing extremely well. Its high precision through-the-lens metering and sophisticated control systems have proven their value, and the system is currently being used to test a silicon-carbide based ladder with a better precision than was achievable with the previous LCFI apparatus. Some preliminary results from this study are shown in Figure 31: it is clear that this ladder exhibits no pathological behaviour under temperature cycling, acquiring a moderate (and reproducible) extra bow. It should be noted that the overall bow present even at room temperature is due to the techniques used to cut the silicon carbide foam. As this technology looks promising, future foam samples will be ordered to precise specifications for prototype ladder building. Note that the "spike" at about 150 mm is caused by the butt joint between the two pieces of 25 μm thick silicon.

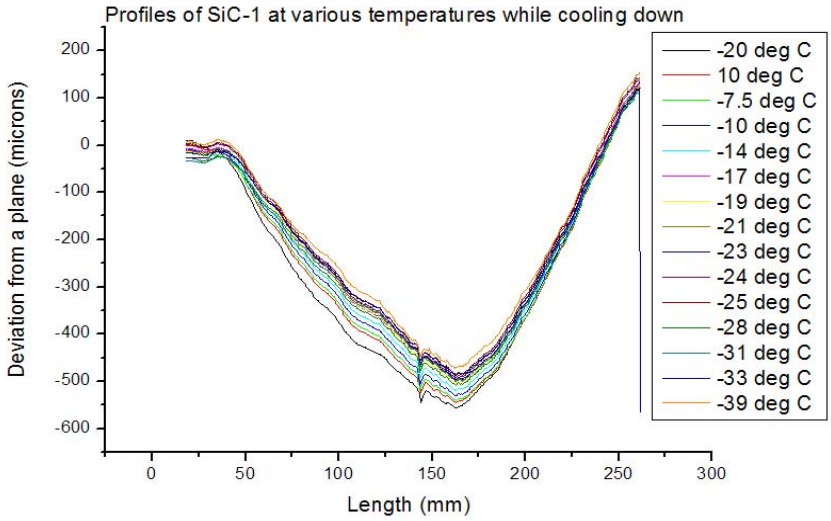

Figure 31 Results of surveying of SiC ladder using new laser apparatus.

Cooling studies

Cooling studies are currently concentrating on the tuning of a Computational Fluid Dynamics model to match the results from a test rig. The CFD model (some results from which are shown in Figure 32) has been extended to include modelling of the solid domain, as well as the fluid itself. This allows the effects of conduction through the structure to be calculated.

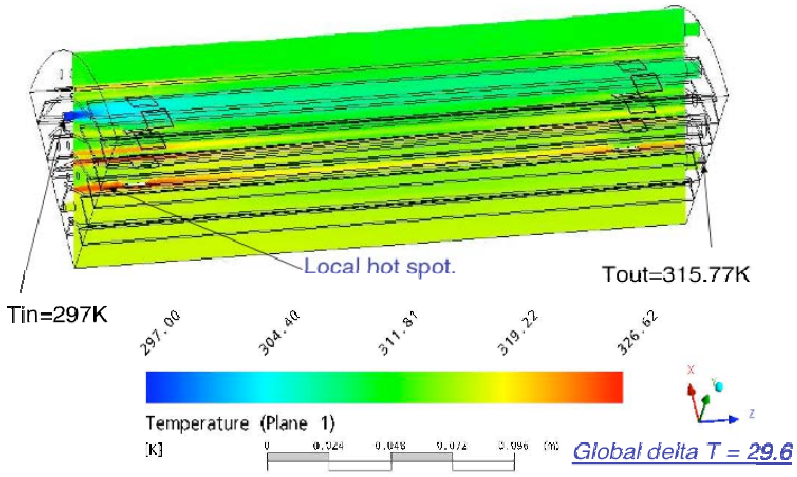

Figure 32 Results obtained using Computational Fluid Dynamics model extended to include properties of solids in the VXD.

Silicon stress studies

LCFI has developed a simple model of the stresses that develop when silicon is processed to produce CCDs. The goal is to provide a description of bowing in thinned silicon sensors that can be used in FEA simulations of potential mechanical structures to evaluate their performance, and to gain a basic understanding of the mechanisms that are likely to be the main factors causing the bowing to occur.

The model developed consists of a thin homogenous layer of material lying over 25 μm of silicon. A temperature change can be applied to force the silicon to curl. It was shown that the observed degree of curling in previously studied samples is consistent with a layer of approximately one micron-thick thermal silicon dioxide, as would be deposited during gate oxidation for example. This effect is also reported in the literature[14], and it is interesting that the differential strain in the SiO_2 layer is almost entirely from the differential thermal contraction. The oxide is observed empirically to relax at high temperatures, so removing any stress from the volume expansion during formation, even though the effect that reduces its viscosity sufficiently to allow "creep" is not understood.

Future plans

The construction and testing of ladders using RVC and SiC foams will continue, greatly aided by the new glue robot and surveying equipment described above. Improved procedures for handling the foams will be developed. Investigations of ladder-based VXDs will continue with studies of possible designs for the bulkheads needed to support the sensor ladders.

An alternative approach to VXD construction is to mount the sensors directly on a shell structure, rather than on individual ladders. First ideas have been discussed and a possible carbon fibre shell structure will be designed.

The comparisons of the Fluid Dynamics calculations and model measurements will be completed.

Summary

As outlined in this document, the Linear Collider Flavour Identification Collaboration has made significant progress in the period since its last report to the DESY Physics Research Committee. Development of the Column Parallel Charge-Coupled Device (CPCCD) has progressed to the extent that the Collaboration is now able to readout the high speed versions of its CPC2, the latest CPCCD design, at frequencies of 45 Mhz, close to the 50 MHz needed for their operation in the inner layer of the Vertex Detector (VXD) of the International Linear Collider (ILC).

A new Column Parallel Readout (CPR) chip has been developed for the CPC2, the CPR2. This adds digital functionality such as threshold application, cluster finding and data sparsification to the amplification and digitisation circuitry that was present on the previous readout chip.

A further major step towards demonstrating that the CPCCD is a viable technology for the VXD has been the successful bump-bonding of the CPC2 to the CPR2 and first successful operation of CPC2/CPR2 assemblies.

The Collaboration has also produced its first Column Parallel Drive chip, the CPD1, which is designed to provide the high current and high frequency drive signals necessary for the CPCCDs. Results of the first tests of this device indicate that it will indeed be able to provide both the currents and frequencies needed to drive the CPCCDs in the VXD.

All the above developments require ancillary electronics systems for their testing, and a continuing programme of development of motherboards has been pursued to allow the new sensors, readout and drive chips to be tested.

Investigations of support structures for the sensors for the VXD have continued and two carbon foam materials identified, reticulated vitreous carbon and silicon carbide, which first test indicate will allow the construction of ladders with a thickness of about $0.1\%X_0$. Test ladders have been manufactured using these materials, with the help of a recently commissioned glue robot, and surveyed using laser surveying equipment of improved precision.

The LCFI Vertex Package will be released at the Linear Collider Workshop at DESY. This sophisticated vertex reconstruction software, with its associated flavour and quark charge identification algorithms, will allow both LCFI and other groups to investigate the physics made accessible at the ILC by the VXD including all the experimental effects simulated in the Geant4-based Monte Carlo Mokka and the influence of the reconstruction algorithms implemented in the reconstruction package MarlinReco. Further, the package is designed to be compatible with other ILC reconstruction and simulation frameworks.

References

[1] http://www.desy.de/f/prc/docs_rd/prc_rd_01_01_update_02_05.pdf
[2] http://lcio.desy.de/
[3] D. Jackson, NIM **A 388** (1997) 247.
[4] N. deGroot, 'From Pixels to Physics', contribution to the 10th International Workshop on Vertex Detectors, VERTEX 2001, Brunnen, Switzerland.
[5] R. Hawkings, 'Vertex detector and flavour tagging studies for the TESLA linear collider', LC-PHSM-2000-021.
[6] http://forum.linearcollider.org
[7] N. Nethercote, J. Seward, 'Valgrind: A Program Supervision Framework', Electronic Notes in Theoretical Computer Science 89 No.2, 2003; web site at http://valgrind.org/
[8] http://www.doxygen.org/
[9] D. Jackson, NIM **A 388** (1997) 247.
[10] A. Raspereza *et al.*, 'Tracking in the LDC Detector, including VXD, FTD, SIT and TPC', contribution to the ECFA and GDE ILC Workshop, Valencia, November 2006.
[11] K. Desch, Th. Kuhl, 'Simulation of Hadronic Branching ratios of a Standard Model like light Higgs Boson at TESLA at 350 GeV', LC-note in preparation (private communication).
[12] J.L. Hewett, 'Indirect Collider Signals for Extra Dimensions', Physical Review Letters **82** (1999) 4765; S. Riemann, 'Fermion-Pair Production at a Linear Collider - A Sensitive Tool for New Physics Searches', LC-TH-2001-007.
[13] M. Battaglia *et al.*, 'Physics Benchmarks for the ILC Detectors', report of the ILC Benchmark panel, hep-ex/0603010.
[14] S.M. Hu, J. Appl. Phys. 70 (6) 1991.